

Computational thresholds for the fixed-magnetization Ising model

Charlie Carlson* Ewan Davies† Alexandra Kolla‡ Will Perkins§

November 5, 2021

Abstract

The ferromagnetic Ising model is a model of a magnetic material and a central topic in statistical physics. It also plays a starring role in the algorithmic study of approximate counting: approximating the partition function of the ferromagnetic Ising model with uniform external field is tractable at all temperatures and on all graphs, due to the randomized algorithm of Jerrum and Sinclair.

Here we show that hidden inside the model are hard computational problems. For the class of bounded-degree graphs we find computational thresholds for the approximate counting and sampling problems for the ferromagnetic Ising model at fixed magnetization (that is, fixing the number of $+1$ and -1 spins).

In particular, letting $\beta_c(\Delta)$ denote the critical inverse temperature of the zero-field Ising model on the infinite Δ -regular tree, and $\eta_{\Delta,\beta,1}^+$ denote the mean magnetization of the zero-field $+$ measure on the infinite Δ -regular tree at inverse temperature β , we prove, for the class of graphs of maximum degree Δ :

1. For $\beta < \beta_c(\Delta)$ there is an FPRAS and efficient sampling scheme for the fixed-magnetization Ising model for all magnetizations η .
2. For $\beta > \beta_c(\Delta)$, there is an FPRAS and efficient sampling scheme for the fixed-magnetization Ising model for magnetizations η such that $|\eta| > \eta_{\Delta,\beta,1}^+$.
3. For $\beta > \beta_c(\Delta)$, there is no FPRAS for the fixed-magnetization Ising model for magnetizations η such that $|\eta| < \eta_{\Delta,\beta,1}^+$ unless $\text{NP}=\text{RP}$.

*Department of Computer Science, University of Colorado Boulder, chca0914@colorado.edu.

†Department of Computer Science, University of Colorado Boulder, ewan.davies@colorado.edu.

‡Department of Computer Science, University of Colorado Boulder, Department of Computer Science and Engineering, University of California Santa Cruz, alexandra.kolla@colorado.edu. Supported in part by NSF grant CCF-1452923

§Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago, math@willperkins.org. Supported in part by NSF grants DMS-1847451 and CCF-1934915.

1 Introduction

The Ising model is a mathematical model of a magnetic material, fundamental in the study of phase transitions in statistical physics. The Ising model is a probability distribution over cuts in a graph, and its partition function is the weighted sum over all cuts in the graph, connecting the physics of the model to combinatorial structures in computer science. In the field of approximate counting in computer science, the ferromagnetic Ising model plays a special role along with the monomer-dimer model as models for which approximating the partition function is tractable on all graphs and at all temperatures [JS89, JS93].

Conditioning on the magnetization of the model corresponds to fixing the balance of the random cut generated. In particular, at zero magnetization (an equal number of plus and minus spins), the Ising model is a probability distribution on bisections of the graph. In the study of spin models on sparse random graphs in physics, it has long been known that conditioning on zero magnetization can turn a ferromagnetic system into a glassy system [MP87] (i.e. fixing the magnetization can drastically change the model and induce slow dynamics). This suggests that lurking inside the tractable computational problems associated to the Ising model there may be hard problems accessible by fixing the magnetization.

We make this idea concrete in a complexity-theoretic sense by reducing NP-hard balanced cut problems to approximating the partition function of the Ising model at fixed magnetization. Specifically we find computational thresholds for approximate counting and sampling in the ferromagnetic Ising model at fixed magnetization on bounded degree graphs. When the inverse temperature β is small (smaller than the critical β on the infinite Δ -regular tree) there are efficient algorithms at all magnetizations. When β is large (larger than the critical β) then there is a computational threshold: for magnetizations η small in absolute value the computational problems are hard; for η large in absolute value the problems are tractable.

We first define the Ising model and the relevant properties of the model on the infinite regular tree, then state our main results.

1.1 The Ising model on graphs and trees

The Ising model on a finite graph G at inverse temperature β and activity λ is the probability distribution $\mu_{G,\beta,\lambda}$ on $\Sigma_G := \{\pm 1\}^{V(G)}$ defined by

$$\mu_{G,\beta,\lambda}(\sigma) = \frac{e^{\frac{\beta}{2} \sum_{(u,v) \in E(G)} \sigma_u \sigma_v} \lambda^{M(\sigma)}}{Z_G(\beta, \lambda)}$$

where $M(\sigma) = \sum_{v \in V(G)} \sigma_v$ and

$$Z_G(\beta, \lambda) = \sum_{\sigma \in \Sigma_G} e^{\frac{\beta}{2} \sum_{(u,v) \in E(G)} \sigma_u \sigma_v} \lambda^{M(\sigma)}.$$

The probability distribution $\mu_{G,\beta,\lambda}$ is the Gibbs measure and $Z_G(\beta, \lambda)$ is the partition function. When $\beta \geq 0$ the model is ferromagnetic, and we will always assume this in what follows. In statistical physics the activity is often written as $\lambda = e^h$ where h is the external field, and so we will call the unbiased case $\lambda = 1$ the zero-field Ising model.

The quantity $M(\sigma)$ is the *magnetization* of the configuration σ . The normalized *mean magnetization* of the Ising model is

$$\eta_G(\beta, \lambda) = \frac{\langle M(\sigma) \rangle_{G, \beta, \lambda}}{|V(G)|},$$

where $\langle \cdot \rangle_{G, \beta, \lambda}$ denotes expectation with respect to the Ising model.

We can also define the Ising model at fixed magnetization. For $k \equiv |V(G)| \bmod 2$, $|k| \leq |V(G)|$, let $\Sigma_G(k) = \{\sigma \in \Sigma_G : M(\sigma) = k\}$ be the subset of Ising configurations with magnetization k . Then the Ising model on G at inverse temperature β and fixed magnetization k is the distribution $\nu_{G, \beta, k}$ on $\Sigma_G(k)$ defined by

$$\nu_{G, \beta, k}(\sigma) = \frac{e^{\frac{\beta}{2} \sum_{(u, v) \in E(G)} \sigma_u \sigma_v}}{Z_G^{\text{fix}}(\beta, k)}$$

where

$$Z_G^{\text{fix}}(\beta, k) = \sum_{\sigma \in \Sigma_G(k)} e^{\frac{\beta}{2} \sum_{(u, v) \in E(G)} \sigma_u \sigma_v}.$$

The distribution $\nu_{G, \beta, k}$ is simply the Ising model at inverse temperature β (and arbitrary activity $\lambda > 0$) conditioned on the event $\sigma \in \Sigma_G(k)$. The fixed-magnetization partition function $Z_G^{\text{fix}}(\beta, k)$ is the coefficient of λ^k when interpreting $Z_G(\beta, \lambda)$ as a Laurent polynomial in λ .

The Ising model can be defined on the infinite Δ -regular tree $\mathbb{T}_\Delta$ via the DLR equations [Dob68, LR69] or as a weak limit of Ising models on finite-depth trees with given boundary conditions. Infinite regular trees are important in computer science as ‘optimal’ expanders, and here we use a known relationship between the Ising model on random regular graphs and on the infinite tree. Depending on the parameters β, λ there may be a unique infinite-volume Gibbs measure on $\mathbb{T}_\Delta$ or there may be multiple measures. The critical inverse temperature is $\beta_c(\Delta) = \log \frac{\Delta}{\Delta-2}$: for $\beta < \beta_c$ there is a unique Gibbs measure for all λ and for $\beta > \beta_c$ there can be multiple measures if λ is close enough to 1 [Lyo89]. We will be interested in one particular Gibbs measure on $\mathbb{T}_\Delta$, the ‘+’ measure induced by the weak limit of finite-depth trees with the all + boundary conditions. We denote this measure $\mu_{\Delta, \beta, \lambda}^+$. By the FKG inequality $\mu_{\Delta, \beta, \lambda}^+$ stochastically dominates all other Gibbs measures on $\mathbb{T}_\Delta$ with the same parameters.

We let $\eta_{\Delta, \beta, \lambda}^+$ denote the expected value of the spin at the root of $\mathbb{T}_\Delta$ under $\mu_{\Delta, \beta, \lambda}^+$ (equivalently, the expected value of the spin at any fixed vertex since $\mu_{\Delta, \beta, \lambda}^+$ is translation invariant). Then the magnetization of the measure $\mu_{\Delta, \beta, \lambda}^+$ is

$$\eta_{\Delta, \beta, \lambda}^+ = \tanh \left(L^* + \operatorname{artanh}(\tanh L^* \tanh \frac{\beta}{2}) \right),$$

where L^* is the largest solution to

$$L^* = \log \lambda + (\Delta - 1) \operatorname{artanh}(\tanh L^* \tanh \frac{\beta}{2}).$$

See Section 4 for more details and a derivation.

The phase transition on $\mathbb{T}_\Delta$ manifests itself via the following ‘spontaneous magnetization’ phenomenon [Lyo89]:

1. For $\beta < \beta_c(\Delta)$, $\eta_{\Delta,\beta,1}^+ = 0$.
2. For $\beta > \beta_c(\Delta)$, $\eta_{\Delta,\beta,1}^+ > 0$.

1.2 Computational problems and computational thresholds

There are two main computational problems associated to a spin model like the Ising model. The approximate counting problem asks for an ε -relative approximation to the partition function Z_G ; that is, a number $\hat{Z}$ so that $(1 - \varepsilon)Z_G \leq \hat{Z} \leq (1 + \varepsilon)Z_G$. An FPTAS is an algorithm that provides such an approximation and runs in time polynomial in $|V(G)|$ and $1/\varepsilon$. An FPRAS is a randomized algorithm that provides such an approximation with probability at least $2/3$ and runs in time polynomial in $|V(G)|$ and $1/\varepsilon$. The approximate sampling problem is to output a sample σ with distribution $\hat{\mu}$ so that $\|\mu_G - \hat{\mu}\|_{TV} < \varepsilon$. An efficient sampling scheme is a randomized algorithm that satisfies this guarantee and runs in time polynomial in $|V(G)|$ and $\log(1/\varepsilon)$ ¹.

Jerrum and Sinclair gave an FPRAS for the ferromagnetic Ising model for all graphs, all inverse temperatures β , and all choices of the activity λ [JS93]². Via self-reducibility of the random cluster representation of the Ising model, this gives an efficient sampling scheme as well [RW99].

On the other hand, for the anti-ferromagnetic Ising model (and the hard-core model of weighted independent sets), the approximate counting and sampling problems are NP-hard in general, and for the class of bounded degree graphs precise computational thresholds are known. The results of Weitz [Wei06], Sly [Sly10], Sly–Sun [SS14], Galanis–Štefankovič–Vigoda [GSV16], and Sinclair–Srivastava–Thurley [SST14] show that for these models (and for β large enough in the case of the anti-ferromagnetic Ising model) there is a computational threshold at some critical activity $\lambda_c = \lambda_c(\Delta, \beta)$. In the case of the hard-core model, there is an FPTAS for $Z_G(\lambda)$ for $\lambda < \lambda_c$ and graphs G of maximum degree Δ while for $\lambda > \lambda_c(\Delta)$ there is no FPRAS unless NP=RP.

For the ferromagnetic Ising model there are no such computational thresholds. But one can ask instead for approximation algorithms for coefficients of the partition function or approximate sampling algorithms for the Ising model at fixed magnetization.

For $\beta \geq 0$ and $\eta \in [-1, 1]$, let $\text{FIXED-ISING}(G, \beta, \eta)$ be the problem of computing the partition function $Z_G^{\text{fix}}(\beta, k)$ of the n -vertex graph G , where k is the largest integer such that $k \equiv n \pmod{2}$ and $k \leq \eta n$. In other words, $k = 2\lfloor(\eta + 1)n/2\rfloor - n$. The associated sampling problem is to sample spin assignments from the measure $\nu_{G,\beta,k}$. The restriction on the parity of k is simply to ensure that configurations of magnetization k exist. Abusing notation slightly we will refer to both η and k as the magnetization, but it will be clear from context what is meant.

This is the setting of the Kawasaki dynamics for the Ising model [Kaw66, Kaw72]: a

¹Sometimes the required dependence of the running time for an efficient approximate sampler is taken to be polynomial in $1/\varepsilon$ instead of $\log(1/\varepsilon)$; we use the stronger definition here.

²In fact the algorithm works in the case of non-uniform activities, as long as they are consistent: all at least 1 or all at most 1. The general case of approximating the partition function with non-uniform activities is #BIS-hard [GJ07].

conservative dynamics with stationary distribution $\nu_{G,\beta,k}$ that at each step proposes a swap of nearest-neighbor spins. Understanding the convergence properties of the Kawasaki dynamics on subsets of $\mathbb{Z}^d$ is a deep mathematical problem [LY93, Yau96, CCM99, CM00]. In this paper we address the problem on general graphs from the perspective of computational complexity.

1.3 Our results

In what follows we always assume $\beta \geq 0$ and $\Delta \geq 3$. When $\beta < \beta_c(\Delta)$ we give efficient approximate counting and sampling algorithms for all magnetizations.

Theorem 1. *Let $\Delta \geq 3$ and $\beta < \beta_c(\Delta)$. Then for all $\eta \in [-1, 1]$ there is an FPRAS and efficient sampling scheme for $\text{FIXED-ISING}(G, \beta, \eta)$ for graphs of maximum degree Δ .*

Theorem 1 can be deduced fairly easily from known results, essentially following the framework of [DP21]. To sample from configurations with a given magnetization, we follow the standard approach of finding a suitable activity parameter for the Gibbs measure $\mu_{G,\beta,\lambda}$ so that the probability of hitting the desired magnetization is not too small (at least inverse polynomial), and then sampling from the Ising model, rejecting samples until we obtain one with the correct magnetization. Because efficient sampling algorithms for the Ising model exist for all β, λ this approach works provided that a suitable activity parameter exists. By continuity, there is an activity that gives the correct mean magnetization, and because the partition function (as a function of λ) is uniformly zero-free in a sector in the complex plane [PR20], the magnetization obeys a central limit theorem [MS19], giving the required inverse polynomial lower bound.

The main results of the paper are for the supercritical case, $\beta > \beta_c(\Delta)$. Here we prove that there is a computational threshold at an explicit $\eta_c = \eta_c(\Delta, \beta) \in (0, 1)$ so that approximation is hard for $|\eta| < \eta_c$ but tractable for $|\eta| > \eta_c$. In fact, $\eta_c(\Delta, \beta) = \eta_{\Delta,\beta,1}^+$, the mean magnetization of the zero-field + measure on $\mathbb{T}_\Delta$.

Theorem 2. *Let $\Delta \geq 3$, $\beta > \beta_c(\Delta)$, and $\eta_c = \eta_{\Delta,\beta,1}^+$.*

- (a) *For all η with $|\eta| > \eta_c$ there is an FPRAS and efficient sampling scheme for $\text{FIXED-ISING}(G, \beta, \eta)$ for graphs of maximum degree Δ .*
- (b) *Unless $\text{NP}=\text{RP}$, for all η with $|\eta| < \eta_c$ there is no FPRAS for $\text{FIXED-ISING}(G, \beta, \eta)$ for graphs of maximum degree Δ .*

In (b) our proof in fact shows that given Δ, β there is some $\zeta > 0$ such that unless $\text{NP}=\text{RP}$, there is no polynomial-time algorithm for $\text{FIXED-ISING}(G, \beta, \eta)$ which achieves a multiplicative approximation of e^{n^ζ} on n -vertex graphs G of maximum degree Δ .

The infinite regular tree plays several roles in the proof of Theorem 2. For the hardness results, non-uniqueness for the zero-field Ising model on the tree at $\beta > \beta_c$ corresponds to ‘phase coexistence’ of the model on the random Δ -regular graph [DM10]. Phase coexistence allows us to use random graphs as gadgets, as Sly does in establishing a computational threshold for the hard-core model [Sly10] (and as is done in subsequent hardness proofs, e.g. [SS14, GSV15, CGG⁺16]). Our analysis of the hardness reduction requires new techniques

to account for the fixed-magnetization constraint; we give an overview of the approach in the next section.

For the algorithmic results, the $+$ measure on the infinite regular tree is the solution to a problem from extremal graph theory that is essential for the proof of Theorem 2.

For the ferromagnetic Ising model with activity $\lambda > 1$, what is the maximum mean magnetization over all graphs of maximum degree Δ ? We prove that the magnetization of the $+$ measure on the infinite Δ -regular tree is an upper bound, and this value is approached by that of the random Δ -regular graph in the $n \rightarrow \infty$ limit. The following result is the main combinatorial result of our paper.

Theorem 3. *For all graphs G of maximum degree Δ , all $\lambda \geq 1$, and all $\beta \geq 0$,*

$$\eta_G(\beta, \lambda) \leq \eta_{\Delta, \beta, \lambda}^+.$$

By integrating the mean magnetization from $\lambda = 1$ to ∞ , this theorem implies the Δ -regular case of a result of Ruozzi which states that the ‘Bethe approximation’ is a lower bound on the normalized partition function of the ferromagnetic Ising model [Ruo12]. In combinatorics, results of this type belong to the field of extremal problems for bounded-degree graphs: maximizing or minimizing observables of statistical physics models over given classes of graphs, like the occupancy fraction of the hard-core or monomer-dimer models [DJPR17]. The area is surveyed by Zhao in [Zha17] and Csikvári describes several cases in which the optimal bound on a partition function is given by an analogous quantity on an infinite regular tree [Csi16]. Bounds on observables such as the mean magnetization or occupancy fraction are stronger than bounds on the partition function, and to the best of our knowledge Theorem 3 is the first case in which the infinite tree is proved to be extremal for an observable.

Theorem 3 implies the following extremal spontaneous magnetization result, which is what we use to guarantee the effectiveness of our algorithm. Define

$$\eta^*(\Delta, \beta) = \lim_{\lambda \rightarrow 1^+} \sup_{G \in \mathcal{G}_\Delta} \eta_G(\beta, \lambda),$$

where $\mathcal{G}_\Delta$ is the class of graphs of maximum degree Δ . Then $\eta^*(\Delta, \beta) = \eta_{\Delta, \beta, 1}^+$. The lower bound comes from taking a sequence of random Δ -regular graphs, while the upper bound follows from Theorem 3. We describe in the next section the content of our algorithmic results for $\beta > \beta_c$: that $\eta_c(\Delta, \beta) = \eta^*(\Delta, \beta) = \eta_{\Delta, \beta, 1}^+$.

1.4 Overview of the techniques

1.4.1 Algorithms

For the algorithmic results of Theorem 2, we aim to apply the same type of algorithm as in Theorem 1: find an activity λ so the mean magnetization is close to the target magnetization, and prove that the probability of hitting the mean is not too small. Again by continuity, there is an activity with the correct mean magnetization, but the distribution may not be concentrated around its mean. For instance, taking $\lambda = 1$ gives 0 mean magnetization by symmetry, but if $\beta > \beta_c$, then hitting 0 magnetization on the random regular graph

is exponentially unlikely. So our question becomes: given an arbitrary graph of maximum degree Δ and a desired magnetization η , is the magnetization under $\mu_{G,\beta,\lambda}$ guaranteed to be concentrated around its mean when λ is chosen so that the mean magnetization is (close to) η ? The answer to this question is given by the Lee–Yang theorem [LY52] in combination with Theorem 3, which guarantees that to achieve a mean magnetization $\eta > \eta^*(\Delta, \beta)$ we can pick an activity λ bounded away from 1 independent of n . The Lee–Yang theorem then gives the zero-freeness result that provides us with the required central limit theorem.

Our proof of Theorem 3 is an extension of an approach used by Krinsky [Kri75] to prove the result for infinite lattices like $\mathbb{Z}^d$ (or more generally graphs satisfying vertex and edge transitivity). The theorem (and the paper [Kri75] that inspired it) may be of independent interest in combinatorics and algorithms. The proof of Theorem 3 relies heavily on correlation inequalities, namely the GKS inequalities [Gri67, KS68], and identities due to Thompson [Tho71]. The techniques are distinct from previous approaches in this area of extremal graph theory such as the entropy method [Kah01], occupancy method [DJPR17], and inductive approaches [Csi17, SSSZ20].

1.4.2 Hardness

To prove a matching hardness result, we must overcome the barrier of the tractability of approximating the Ising partition function. This rules out the approach used in [DP21] for proving hardness of approximating the number of independent sets of a given size, namely reducing approximating the partition function to approximating a fixed coefficient of the partition function. Instead, we use the fact that imposing the fixed-magnetization constraint fundamentally alters the behavior of the model. When highly connected components of a graph are connected with a relatively sparse set of edges, the fixed-magnetization, zero-field ferromagnetic Ising model exhibits a kind of global anti-ferromagnetic behavior due to the constraint on the magnetization: the spins on each highly connected component will align, but the number of components that pick each spin will be essentially determined by the constraint. This behavior is what allows us to prove hardness. We use a probabilistic analysis of the fixed-magnetization Ising model to show that a gadget construction based on that of [Sly10] can be used to reduce an NP-hard cut problem to approximating the fixed-magnetization Ising partition function. To illustrate our methods we sketch a simplified version of the proof for zero magnetization.

Similar to previous approaches, our gadget G is essentially a random Δ -regular bipartite graph with some edges removed and trees attached to create ‘terminal vertices’ of degree $\Delta - 1$. Given an instance H of MIN-BISECTION, we replace each vertex of H by a copy of the gadget G and then join a number of terminal vertices of the appropriate copies of the gadget graph for each edge of H . When $\beta > \beta_c$, the Ising model on a single gadget G exhibits *phase coexistence*, with a bimodal distribution of either many more $+$ spins than $-$ spins or vice versa. The phase coexistence property of each gadget is so strong that when we take the collection of gadgets joined by the crossing edges and condition the Ising model on zero magnetization, the phase coexistence property on each gadget persists, and zero-magnetization is achieved (with high probability) by having an equal number of gadgets in each phase. Showing this involves proving a local central limit theorem and large deviation results for the magnetization of a collection of gadgets conditioned on an arbitrary spin assignment to the

set of terminal vertices. This shows that the dominant contribution to the zero-magnetization partition function is given by configurations whose gadget phase assignments encode minimum bisections of H , and this in turn implies that a good approximation algorithm for the partition function can recover a minimum bisection.

The proof of the local central limit theorem conditioned on the phases of the gadgets is a new technical ingredient in our proof. It involves bounding the moments of the magnetization on a single gadget, conditioned on a phase, and employing a Fourier analytic proof of a local central limit theorem.

The full proof and the general case of $\eta \neq 0$ are only slightly more complex. Broadly, the same approach works except we reduce from a generalization of MIN-BISECTION, γ -MIN-EXACT-BALANCED-CUT (γ -MEBC), that requires the partition of a vertex set of size N to have part sizes $\lfloor \gamma N \rfloor$ and $\lceil (1 - \gamma)N \rceil$. It is convenient to add to the collection of gadget graphs some isolated vertices which smooth out certain parts of the analysis. In particular, it helps in proving the local central limit theorem. We choose γ as a function of Δ , β , and η , and we prove that when the Ising model on the collection of gadget graphs is conditioned to have magnetization η , with high probability the phases of the gadgets are split in fractions γ and $1 - \gamma$. Then a good approximation algorithm for the η -magnetization partition function can recover a minimum γ -balanced cut.

1.5 Related work

The algorithmic problem of sampling configurations of a fixed magnetization (or fixed size, in the case of independent sets) is the problem of sampling from the ‘canonical ensemble’ in the language of statistical physics (in contrast to the ‘grand canonical ensemble’ of the usual Ising or hard-core model). Work on this problem goes back to the very first Markov Chain Monte Carlo algorithm designed to sample from the canonical ensemble of hard spheres [MRR⁺53]. Conservative dynamics such as these are still among the most used in current scientific applications (e.g., [BKW09]). Grand canonical ensembles are generally more amenable to mathematical analysis due to their conditional independence properties, and much is known about both specific algorithms for sampling from these distributions (e.g. Glauber dynamics [MS13], random-cluster dynamics [GJ18]) and about the computational complexity of the approximate counting and sampling problems for these models.

The computational complexity of approximately counting and sampling independent sets of a given size in bounded-degree graphs was recently addressed by Davies and Perkins who proved a computation threshold for these problems [DP21]. As in Theorem 2, the threshold is given in terms of an extremal graph theory problem: that of minimizing the occupancy fraction over $G \in \mathcal{G}_\Delta$. Faster algorithms and an FPTAS up to the threshold for this problem were recently given in [JPSS21].

The use of random graphs as gadgets in hardness reductions was pioneered by Dyer, Frieze, and Jerrum [DFJ02] and used by Sly in identifying the computational threshold for the hard-core model [Sly10], with further applications in [SS14, GSV15, CGG⁺16, GSVY16] among others. In particular, a detailed understanding of the moments of the partition function Z_G for random regular graphs is now known, and, via the small subgraph conditioning method, concentration results for Z_G . We use this understanding extensively in Section 3.

Finally, the Ising model at fixed magnetization has been studied extensively in both mathematics and physics, on $\mathbb{Z}^d$ and on random graphs [MP87]. Conditioning the ferromagnetic Ising model on zero magnetization has the effect of introducing ‘frustration’: the impossibility of satisfying all edge constraints simultaneously.

At zero temperature ($\beta = \infty$), the zero-magnetization Ising model is simply the uniform distribution on min-bisections of a graph; finding the size of the min bisection has long been known to be NP-hard [GJS74]. The min-bisection problem is also studied on random graphs from the perspective of statistical physics [PIG⁺08, ZB10, DSW04, DMS17]. Our work is an exploration of the worst-case computational complexity of the positive temperature regime of this problem.

1.6 Questions and future directions

Though we do not pursue it in this extended abstract, it is likely that the techniques of Jain, Perkins, Sah, and Sawhney [JPSS21] can be used to improve the algorithmic results of Theorems 1 and 2 in two ways:

1. Obtain an FPTAS (efficient deterministic approximation algorithm) for $\text{FIXED-ISING}(G, \beta, \eta)$ for the same range of parameters for which we obtain an FPRAS.
2. Improve the running time of our approximate sampling algorithm to $\tilde{O}(n \log n)$.

We have shown here a computational threshold for the fixed-magnetization Ising model. One can also ask what is achievable with a specific algorithm widely used in scientific applications, namely the Kawasaki dynamics. We conjecture that the Kawasaki dynamics mix rapidly on all graphs of maximum degree Δ for the same set of parameters for which we provide an FPRAS. In fact there are two versions of the Kawasaki dynamics: the *local flip* dynamics in which at each step a swap of spins across an edge is proposed; and the *global flip* dynamics in which at each step a swap of arbitrary spins in the graph is proposed. We conjecture that both versions mix in polynomial time for the parameters above; we further conjecture that the global flip dynamics mix in time $O(n \log n)$.

Conjecture 1. *For $\beta < \beta_c(\Delta)$, the Kawasaki dynamics mix in time polynomial in n for any fixed magnetization and any graph G of maximum degree Δ on n vertices.*

For $\beta > \beta_c(\Delta)$ and $|\eta| > \eta_c(\Delta, \beta)$ the Kawasaki dynamics mix in time polynomial in n for any fixed magnetization $k \geq \eta n$ and any graph G of maximum degree Δ on n vertices.

For the global flip dynamics, the mixing time in both cases is $O(n \log n)$.

In the previous uses of random (bipartite) graphs as gadgets in hardness reductions for approximate counting problems, the gadgets themselves are not in general hard instances for the given problems. In particular, recent results [JKP20, HJP20, CGŠVar, JPPar] show that for parameters sufficiently deep in the given non-uniqueness regimes, random regular graphs are tractable instances for approximate counting and sampling. We ask whether for random graphs there are efficient algorithms anywhere inside the NP-hardness regime.

Question 1. For $\Delta \geq 3$, $\beta > \beta_c(\Delta)$, is there some $|\eta| < \eta_c(\Delta, \beta)$ so that there exist efficient approximate counting and sampling algorithms for $\text{FIXED-ISING}(G, \beta, \eta)$ for random Δ -regular graphs?

1.7 Organization

In Section 2 we provide some of the results we will use in our algorithms and hardness reductions. In Section 3 we give the hardness reduction. In Section 4 we prove Theorem 3, solving the extremal problem that identifies the limit of our algorithmic approach. In Section 5 we prove the algorithmic results.

2 Preliminaries

Recall that $\mathcal{G}_\Delta$ denotes the class of graphs of maximum degree Δ . We use $\mu_{G,\beta,\lambda}$ to denote the Ising model on G at inverse temperature β and activity λ . We will often drop β from the notation when it remains fixed and we will drop λ from the notation in the case $\lambda = 1$ (so $\mu_G = \mu_{G,\beta,1}$ when β is understood from the context). We use the bracket notation $\langle \cdot \rangle_{G,\beta,\lambda}$ to denote expectations with respect to the Ising model, in part to distinguish these expectations from expectations over random graphs in Section 3. For a graph G , let $\Sigma_G = \{\pm 1\}^{V(G)}$. Slightly abusing notation, for $U \subset V(G)$, let $\Sigma_U = \{\pm 1\}^U$. We let $M(\sigma)$ denote the magnetization of a configuration σ and $X(\sigma)$ denote the number of $+$ spins (so $M(\sigma) = 2X(\sigma) - |V(G)|$). We let $\mathbf{X}$ denote the random variable $X(\sigma)$ when σ is drawn from $\mu_{G,\beta,\lambda}$.

We now collect a number of results that we will use in the proofs that follow. The first results are results on zero-free regions for the Ising model partition function, viewed as a (Laurent) polynomial in λ .

Theorem 4 (Lee–Yang [LY52]). For $\beta \geq 0$, $\lambda \in \mathbb{C}$, and any graph G , $Z_G(\beta, \lambda) = 0$ only if $|\lambda| = 1$.

Theorem 5 (Peters–Regts [PR20]). Let $\Delta \geq 3$ and $\beta \in (0, \beta_c(\Delta))$. Then there exists $\theta = \theta(\beta) \in (0, \pi)$ such that for any $\lambda \in \mathbb{C}$ with $|\arg(\lambda)| < \theta$ and any graph $G \in \mathcal{G}_\Delta$ we have $Z_G(\beta, \lambda) \neq 0$.

By the following general result of Michelen and Sahasrabudhe, these zero-freeness results imply central limit theorems for the magnetization of the ferromagnetic Ising model on graphs in $\mathcal{G}_\Delta$ when $\beta < \beta_c(\Delta)$ or when $\lambda > 1$. We apply this result to a random variable counting the number of $+1$ spins in a sample from the Ising model; its generating function is a scaling of $Z_G(\beta, \lambda)$.

Theorem 6 (Michelen–Sahasrabudhe [MS19]). For $n \geq 1$ let $\mathbf{X}_n$ be a random variable taking values in $\{0, \dots, n\}$ with mean μ_n , standard deviation σ_n , and probability generating function f_n . If the roots ζ of f_n satisfy $|\arg(\zeta)| \geq \delta_n$ and $\sigma_n \delta_n \rightarrow \infty$, then $(X_n - \mu_n)/\sigma_n$ converges in distribution to a standard normal random variable.

A central limit theorem for the magnetization in fact implies a local central limit theorem, following the approach of Dobrushin and Tirozzi [DT77] (for spin models on $\mathbb{Z}^d$) and the

results of [JPSS21] for the hard-core model. Let $\mathbf{X}$ denote the number of +1 spins in a sample from the Ising model.

Proposition 7. *Fix $\lambda > 1$ and $\beta \geq 0$. Then for any graph $G \in \mathcal{G}_\Delta$ on n vertices and any non-negative integer ℓ ,*

$$\mu_{G,\beta,\lambda}(\mathbf{X} = \ell) = \frac{1}{\sqrt{2\pi \text{var}(\mathbf{X})}} \exp \left[-\frac{(\ell - \langle \mathbf{X} \rangle_{G,\beta,\lambda})^2}{2 \text{var}(\mathbf{X})} \right] + o(n^{-1/2}),$$

where $\text{var}(\mathbf{X}) = \langle \mathbf{X}^2 \rangle_{G,\beta,\lambda} - \langle \mathbf{X} \rangle_{G,\beta,\lambda}^2$, and where the implied constant in the error term depend only on Δ, β, λ . The same holds for $\beta < \beta_c$, $\lambda \geq 1$, and any $G \in \mathcal{G}_\Delta$.

Moreover, under the conditions above $\text{var}(\mathbf{X}) = \Theta(n)$ where again the implied constants depend only on Δ, β, λ .

We prove Proposition 7 in Appendix A; the proof of the local central limit theorem is analogous to that of [JPSS21, Theorem 1.5] and the proof of the variance bound is analogous to that of [DP21, Lemma 9] and [JPSS21, Lemma 3.2].

For the hardness results, we reduce an NP-hard cut problem to the problem of approximating the Ising model at fixed magnetization. The γ -MIN-EXACT-BALANCED-CUT (γ -MEBC) problem is the problem of finding the minimum of $|E(S, S^c)|$ over all $S \subset V(G)$, $|S| = \lfloor \alpha n \rfloor$, where $n = |V(G)|$. For stronger inapproximability in Theorem 2(b), we apply an inapproximability result due to Bui and Jones [BJ92], though this is not essential to our method: we can reduce from exactly solving γ -MEBC instead.

Theorem 8 (Bui–Jones [BJ92]). *Let γ be a rational number in $(0, 1)$ and let $\varepsilon > 0$. Then γ -MEBC is NP-hard to approximate within an additive error $n^{2-\varepsilon}$ on n -vertex graphs.*

A key ingredient in the algorithmic results are the efficient approximate counting and sampling algorithms for the ferromagnetic Ising model provided by Jerrum and Sinclair and Randall and Wilson.

Theorem 9 (Jerrum–Sinclair [JS93], Randall–Wilson [RW99]). *For all inverse temperatures β and all activities λ , there is an FPRAS and efficient sampling scheme for the Ising model for all graphs G .*

3 Hardness

3.1 The reduction and its properties

Given $\Delta \geq 3$, $\beta > \beta_c(\Delta)$ and $\eta \in [0, \eta_c)$, our goal is to reduce γ -MIN-EXACT-BALANCED-CUT to approximating a fixed-magnetization Ising partition function, for some rational number $\gamma \in ((1 + \eta/\eta_c)/2, 1)$. (By symmetry we need only consider $\eta \geq 0$).

For the reduction we require a gadget $G = G(\Delta, n, \theta, \psi)$ where $\theta, \psi \in (0, 1/8)$ are constants that can be determined later in terms of Δ, β . The gadget is identical to the constructions in [Sly10, GSVY16], which is a balanced bipartite graph on $n_G = (2 + o(1))n$ vertices. The majority of the vertices have degree Δ , and $m = O(n^\theta)$ vertices on each side of G are

designated *terminal vertices* of degree $\Delta - 1$. We detail the construction of the gadget and state its properties after showing how it is used in the reduction.

Let H be a graph on $h = \lfloor n^{\theta/4}/(\Delta - 1) \rfloor$ vertices, which is the input for γ -MEBC. Given G as above and an integer s , we construct a graph H_s^G of maximum degree Δ on $N := hn_G + s$ vertices as follows:

- We include a copy G^x of G for each vertex $x \in V(H)$.
- We include s isolated vertices.
- For each edge $xy \in E(H)$, we include a matching of size $k = \lfloor n^{3\theta/4} \rfloor$ between the left terminals of G^x and left terminals of G^y and a matching of size k between the right terminals of G^x and right terminals of G^y . We do this in such a way that each terminal is used at most once (which is possible since $kh \leq m$).

For reference, our parameter choices are listed here. The parameters Δ, β are fixed and we can compute η_c from them (to arbitrary precision). The parameter η is fixed and satisfies the conditions of the theorem. From these parameters we compute an arbitrary rational number γ such that

$$\frac{1 + \eta/\eta_c}{2} < \gamma < 1,$$

which is possible because $\eta \in [0, \eta_c)$. Suitable choices of $\theta, \psi \in (0, 1/8)$ can be made in terms of Δ, β (see Lemma [III](#)). We are then given an instance H of γ -MIN-EXACT-BALANCED-CUT on h vertices with h sufficiently large, and we choose an n large enough that $h \leq n^{\theta/4}/(\Delta - 1)$. Let $h_+ = \lfloor \gamma h \rfloor$ and $h_- = \lceil (1 - \gamma)h \rceil$ so that the γ -MIN-EXACT-BALANCED-CUT problem is to find the minimum of $|E(S, S^c)|$ over $S \subset V(H)$ with $|S| = h_+$. We insist that h is large enough that $\min\{h_+, h_-\} \geq 1$. Now let

- $m = (\Delta - 1)^{\lfloor \theta \log_{\Delta-1} n \rfloor} = o(n^{1/8})$,
- $m' = (\Delta - 1)^{\lfloor \theta \log_{\Delta-1} n \rfloor + \lfloor \psi \log_{\Delta-1} n \rfloor} = o(n^{1/4})$,
- $n_G = 2(n + m' + m((\Delta - 1)^{\lfloor \psi \log_{\Delta-1} n \rfloor} - 1)/(\Delta - 2)) = (2 + o(1))n$
- $k = \lfloor n^{3\theta/4} \rfloor$,
- s be a non-negative integer such that

$$|2n(h_+ - h_-)\eta_c - \eta[hn_G + s]| \leq \sqrt{nh}, \quad (1)$$

and $s = \Theta(nh)$,

- $N = hn_G + s$,
- $M^* = h_+ - h_-$,
- $\delta > 0$ be small enough as a function of γ and η_c ,
- $\ell = \lfloor N(\eta + 1)/2 \rfloor$.

The parameters m, m', n_G, k relate to the gadget construction that we detail below. There exists an s satisfying (II) with $s = \Theta(nh)$ which can be found in time polynomial in h because our parameter choices mean that (as $h \rightarrow \infty$)

$$2n(h_+ - h_-)\eta_c = (2 + o(1)) \cdot (2\gamma - 1)\eta_c \cdot nh,$$

and

$$\eta[hn_G + s] = (2 + o(1)) \cdot \eta \cdot nh + O(s).$$

Since $(2\gamma - 1)\eta_c > \eta$, for some non-negative integer $s = \Theta(nh)$ the latter can be made within an additive term $O(1)$ of the former (and hence within $\sqrt{nh}$). Finally, N is the number of vertices in the graph H_s^G which we construct in the reduction, M^* is the magnetization of the cuts considered for the γ -MEBC problem on H , and ℓ is such that on H_s^G FIXED-ISING(G, β, η) asks for configurations σ with $X(\sigma) = \ell$ (and the desired fixed magnetization is thus $2\ell - N$).

Throughout this section there are many absolute constants (depending only on Δ, β, η) used and defined. For ease of reading we will make ample use of $O(\cdot)$ and $\Omega(\cdot)$ notation as well as reusing constants C, c etc.

The main result of this section is the following.

Theorem 10. *Given $\varepsilon \in (0, 1)$ there exists $\zeta > 0$ such that there is a randomized, polynomial-time algorithm to construct a graph G as above so that with probability at least $2/3$ the following holds: given an e^{N^ζ} -relative approximation to $Z_{H_s^G}^{\text{fix}}(\beta, 2\ell - N)$ one can compute, in time polynomial in h , an additive $h^{2-\varepsilon}$ approximation to the γ -MIN-EXACT-BALANCED-CUT of H .*

Theorem 10 together with Theorem 8 immediately gives Theorem 2(b).

3.2 The gadget

We use the same gadget construction as in [Sly10, GŠVY16]. The construction is defined by the maximum degree Δ , an integer n and constants $\theta, \psi \in (0, 1/8)$ which then determine the parameters m, m', n_G, k listed above. To construct $G = G(\Delta, n, \theta, \psi)$, let $G' = G'(\Delta, n, \theta, \psi)$ be a random bipartite graph with $n + m'$ vertices on each side obtained by choosing Δ perfect matchings between the sides uniformly at random, and from the final matching removing m' of the edges. With high probability the matchings will be pairwise disjoint sets of edges, so G' is a simple graph. Let U_0 be the set of vertices of degree Δ in G' , and W_0 be the set of vertices of degree $\Delta - 1$.

To form G from G' , on each side partition the m' vertices of degree $\Delta - 1$ into m equal-sized sets, and attach the leaves of a copy of a $(\Delta - 1)$ -ary tree of depth $\lfloor \psi \log_{\Delta-1} n \rfloor$ to each set. Then each side of G' has had m trees each of which contains $O(n^\psi)$ vertices added. The roots of these trees are now the only vertices of degree $\Delta - 1$ in G , and there are m roots that were added to each side. These are the *terminal* vertices which allow us to connect the gadgets together. Let V_0 be the vertex set of G , and R_0 be the terminal vertices.

Constructed in this way, we want to show that various properties of G' and G hold with sufficiently high probability. Many of these properties were verified in [Sly10, GŠVY16], but we require additional control of statistics of the number of $+1$ spins.

Throughout this entire section we fix the inverse temperature $\beta > \beta_c$ and take $\lambda = 1$. Recall that $\mathbf{X}$ denotes the number of $+1$ spins in a sample from the Ising model. We will condition on various *phases* of the Ising model on G and on H_s^G . For G , given an Ising configuration $\sigma \in \Sigma_G$, we say the phase is $+$ if $\sum_{v \in U_0} \sigma_v > 0$ and $-$ if $\sum_{v \in U_0} \sigma_v < 0$. If the sum is 0 then we take the phase to be the spin of some distinguished vertex $u_1 \in U_0$ fixed in advance (arbitrarily). Note that neither the spins of W_0 nor the spins of the trees added to G' in the construction of G appear in the definition of the phase. By symmetry, the probability under μ_G of each phase is exactly $1/2$. We denote the Ising model on G conditioned on the $+$ phase and $-$ phase respectively by $\mu_{G,+}$ and $\mu_{G,-}$. We use $\langle \cdot \rangle_{G,+}$ and $\langle \cdot \rangle_{G,-}$ to denote the corresponding conditional expectation operators. For a spin assignment $\tau \in \Sigma_{R_0}$ to the terminals of G , we include an additional subscript τ to denote conditioning on the event $\{\sigma_{R_0} = \tau\}$ (that is, σ restricted to R_0 is equal to τ).

Let $Z_{G,\alpha}$ be the contribution to the partition function Z_G of the Ising model on the random gadget G from spin assignments in which there are precisely $2\alpha n$ vertices in U_0 of spin $+$. Let $Z_{G,+}$ and $Z_{G,-}$ be the contributions from the $+$ and $-$ phase respectively to Z_G . We have

$$Z_{G,+} = \sum_{1/2 < \alpha \leq 1} Z_{G,\alpha} + \frac{1}{2} Z_{G,1/2},$$

because all contributions with $\alpha > 1/2$ belong to the $+$ phase, but for $\alpha = 1/2$ we break the tie symmetrically so that half the contribution $Z_{G,1/2}$ goes to the $+$ phase. Usually, it suffices to use the upper bound on $Z_{G,\pm}$ obtained by taking the entire contribution from $\alpha = 1/2$ to the phase at hand.

We now state some results from [Sly10, GŠV15, GŠVY16] which are obtained by sophisticated versions of the first and second moment methods and an application of the small subgraph conditioning method [RW94, Jan95]. We will state the results for the $+$ phase, the $-$ phase is completely analogous. Let $\alpha^+ = (1 + \eta_c)/2$, $\alpha^- = (1 - \eta_c)/2$, and

$$q = \frac{(\alpha^+ e^\beta + 1 - \alpha^+)^{\Delta-1}}{(\alpha^+ e^\beta + 1 - \alpha^+)^{\Delta-1} + (\alpha^+ + (1 - \alpha^+) e^\beta)^{\Delta-1}}, \quad (2)$$

which means q is the probability that the root of a $(\Delta - 1)$ -ary tree gets spin $+1$ in the zero-field ‘ $+$ measure’ on the infinite $(\Delta - 1)$ -ary tree (defined analogously to the $+$ measure on the infinite Δ -regular tree). Note that $1/2 < q < \alpha^+$. We use $Q_S^+(\cdot)$ to denote the product measure on Σ_S that assigns probability q to $+1$ spins and probability $1 - q$ to -1 spins, and vice versa for $Q_S^-(\cdot)$.

The following lemma collects previous results on the gadget, most notably the near independence of the terminal spins conditioned on a phase.

Lemma 11 ([Sly10, Proof of Theorem 2.1], [GŠV15, Proof of Lemma B.3], [GŠVY16, Lemma 22]). *Let G be the random graph described above with parameters $\theta, \psi \in (0, 1/8)$. Then there exists $c > 0$ so that for all $\alpha \in [1/2, 1]$,*

$$\frac{\mathbb{E} Z_{G,\alpha}}{\mathbb{E} Z_{G,+}} \leq \frac{C}{\sqrt{n}} e^{-cn(\alpha + \alpha^+)^2},$$

and for all $\alpha \in [0, 1/2]$

$$\frac{\mathbb{E}Z_{G,\alpha}}{\mathbb{E}Z_{G,-}} \leq \frac{C}{\sqrt{n}} e^{-cn(\alpha+\alpha^-)^2}.$$

Moreover, there exist choices of constants $\theta, \psi \in (0, 1/8)$ and $C' > 0$ so that for large enough n , with probability at least $9/10$ over the choice of G , all of the following hold simultaneously:

(i) Conditioned on phase $+$, the terminal spins are approximately independent:

$$\max_{\tau \in \Sigma_{R_0}} \left| \frac{\mu_{G,+}(\sigma_{R_0} = \tau)}{Q_{R_0}^+(\tau)} - 1 \right| \leq n^{-2\theta}.$$

(ii) There exists $\mathcal{B} \subset \Sigma_{W_0}$ so that

- $\mu_{G,+}(\sigma_{W_0} \in \mathcal{B}) \leq \exp(-n^{2\theta})$
- For every $\tau_{W_0} \in \Sigma_{W_0} \setminus \mathcal{B}$,

$$\max_{\tau_{R_0} \in \Sigma_{R_0}} \left| \frac{\mu_{G,+}(\sigma_{R_0} = \tau_{R_0} | \sigma_{W_0} = \tau_{W_0})}{Q_{R_0}^+(\tau_{R_0})} - 1 \right| \leq n^{-3\theta}$$

(iii)

$$Z_{G,+} > \frac{1}{C'} \mathbb{E}Z_{G,+}.$$

The same also hold with $+$ replaced by $-$.

In previous works [Sly10, GSVY16] a version of (iii) with $1/C'$ replaced by a function $o(1)$ as $n \rightarrow \infty$ is used (along with the fact that this weaker bound holds with high probability), but the stated version follows from the small subgraph conditioning method used therein [RW94, Jan95]. In order to handle the fixed-magnetization constraint in our reduction, we show that certain additional properties hold with good probability.

Lemma 12. *For sufficiently large n , with probability at least $8/10$ over the choice of the gadget G described above, the following hold simultaneously, for all choices of $\tau \in \Sigma_{R_0}$ (and for $+$ replaced by $-$ as well):*

(a)

$$\langle |\mathbf{X} - 2n\alpha^+| \rangle_{G,+, \tau} = O(\sqrt{n})$$

(b)

$$\langle |\mathbf{X} - 2n\alpha^+|^2 \rangle_{G,+, \tau} = O(n)$$

(c)

$$\langle |\mathbf{X} - 2n\alpha^+|^3 \rangle_{G,+, \tau} = O(n^{3/2})$$

(d) For $\delta > 0$ as specified above and $t_0 = \delta/(2c_0 h)$ for some constant $c_0 > 1/4$,

$$\langle e^{t_0(\mathbf{X} - 2\alpha^+ n)} \rangle_{G,+, \tau} \leq e^{c_0 t_0^2 n}$$

and

$$\langle e^{t_0(2\alpha^+ n - \mathbf{X})} \rangle_{G,+, \tau} \leq e^{c_0 t_0^2 n}.$$

We prove Lemma 12 in Section 12.

3.3 Proof of Theorem 10

Here we prove Theorem 10 given Lemmas 11 and 12. Let G be a gadget which satisfies the conclusions of these lemmas, and let H_s^G denote the graph constructed from H , G and isolated vertices as above. Let $\hat{H}_s^G$ denote the same graph but without the edges between gadgets (so $\hat{H}_s^G$ consists of h disjoint copies of G and s isolated vertices). We write $\mathcal{E}$ for the set of edges $E(H_s^G) \setminus E(\hat{H}_s^G)$ that lie between gadgets.

Given $\sigma \in \Sigma_{H_s^G}$, let $Y(\sigma) \in \Sigma_H$ be a vector denoting the phases of the gadgets G^x , $x \in V(H)$. We will call such a Y a *phase vector*. For a given $Y \in \Sigma_H$ let $\mu_{H_s^G, Y}$ be the Ising model on H_s^G conditioned on $\{Y(\sigma) = Y\}$. Let $\langle \cdot \rangle_{H_s^G, Y}$ be the corresponding expectation operator. Define $\mu_{\hat{H}_s^G, Y}$ and $\langle \cdot \rangle_{\hat{H}_s^G, Y}$ analogously. For $x \in V(H)$ let R^x be the set of terminals in the gadget G^x , and let R be the union of the terminal vertices in all the copies of the gadget. For a spin assignment $\tau \in \Sigma_R$ to the terminals, we include an additional subscript τ to indicate conditioning on the event that $\{\sigma_R = \tau\}$.

We need two probabilistic results before proving Theorem 10. The first is a large deviation bound for $\mathbf{X}$ conditioned on any phase vector Y and any assignment of terminal spins τ . The second is a local central limit theorem for $\mu_{\hat{H}_s^G, Y, \tau}(\mathbf{X} = \ell)$ when $M(Y) = M^*$, and for an arbitrary terminal spin assignment τ .

Lemma 13. *Assume the gadget G satisfies the conclusions of Lemma 12. Then for any phase vector Y , any $\tau \in \Sigma_W$, with*

$$\nu = \frac{s}{2} + 2n \sum_{x \in V(H)} \alpha^{Y_x},$$

we have

$$\mu_{\hat{H}_s^G, Y, \tau}(|\mathbf{X} - \nu| \geq \delta n) \leq \exp(-\Omega(n/h)),$$

where $\delta > 0$ is a constant defined above.

Proof. Note that to leading order ν is the mean $\langle \mathbf{X} \rangle_{\hat{H}_s^G, Y, \tau}$.

We prove the bound on the upper tail; the proof for the lower tail is identical. Let $\mathbf{X}_s \sim \text{Bin}(s, 1/2)$ be the number of $+$ spins among the s isolated vertices. Then $\langle e^{t_0(\mathbf{X}_s - s/2)} \rangle \leq e^{st_0^2/4} \leq e^{sc_0 t_0^2}$ since $c_0 > 1/4$. Using this along with Lemma 12 we can bound the moment generating function,

$$\langle e^{t_0(\mathbf{X} - \nu)} \rangle_{\hat{H}_s^G, Y, \tau} \leq e^{c_0 t_0^2 n h},$$

for some constant $c > 0$. Then we have

$$\begin{aligned} \mu_{\hat{H}_s^G, Y, \tau}(\mathbf{X} \geq \nu + \delta n) &\leq e^{-t_0 \delta n} \langle e^{t_0(\mathbf{X} - \nu)} \rangle_{\hat{H}_s^G, Y, \tau} \\ &\leq e^{-t_0 \delta n + c t_0^2 n h} = e^{-\frac{\delta n}{4h}} \end{aligned}$$

since $t_0 = \delta/(2ch)$. □

The next lemma is a local central limit theorem for $\mathbf{X}$ with respect to $\mu_{\hat{H}_s^G, Y, \tau}$.

Lemma 14. Assume the gadget G satisfies the conclusions of Lemma 12. Then for any phase vector $Y \in \Sigma_H$ with $M(Y) = M^*$, any $\tau \in \Sigma_W$, and any integer t ,

$$\mu_{\hat{H}_s^G, Y, \tau}(\mathbf{X} = t) = \frac{1}{\sqrt{2\pi\kappa^2}} \exp \left[-\frac{(t - \langle \mathbf{X} \rangle_{\hat{H}_s^G, Y, \tau})^2}{2\kappa^2} \right] + o \left(\frac{1}{\sqrt{nh}} \right),$$

where $\kappa^2 = \text{var}_{\hat{H}_s^G, Y, \tau}(\mathbf{X})$. In particular,

$$\mu_{\hat{H}_s^G, Y, \tau}(\mathbf{X} = \ell) = \Omega \left(\frac{1}{\sqrt{nh}} \right).$$

Proof. We have $s = \Theta(nh)$, and by the independence of disjoint gadgets and the second moment bound in Lemma 12, we have $\kappa^2 = \Theta(nh)$. Moreover, by our choice of s and the fact that $M(Y) = M^*$, we have $|\langle \mathbf{X} \rangle_{\hat{H}_s^G, Y, \tau} - \ell| = O(\sqrt{nh})$, and so the second statement follows from the first.

The proof of the first statement is similar to that of Proposition 7 in Appendix A, but here things are especially simple because of the presence of $s = \Theta(nh)$ isolated vertices.

We start by proving a central limit theorem with the standard method of characteristic functions. Let $\bar{\mathbf{X}} = \frac{1}{\kappa}(\mathbf{X} - \langle \mathbf{X} \rangle_{\hat{H}_s^G, Y, \tau})$ and $\phi_{\bar{\mathbf{X}}}(t) = \langle e^{it\bar{\mathbf{X}}} \rangle_{\hat{H}_s^G, Y, \tau}$. Then

$$\begin{aligned} \phi_{\bar{\mathbf{X}}}(t) &= \left(\frac{1 + e^{it/\kappa}}{2} e^{-t/(2\kappa)} \right)^s \prod_{x \in V(H)} \left\langle e^{\frac{it}{\kappa}(\mathbf{X} - \langle \mathbf{X} \rangle_{G, Y_x, \tau_{Rx}})} \right\rangle_{G, Y_x, \tau_{Rx}} \\ &= \left(1 - \frac{t^2}{8\kappa^2} + O(\kappa^{-3}) \right)^s \prod_{x \in V(H)} \left(1 - \frac{t^2 \text{var}_{G, Y_x, \tau_{Rx}}(\mathbf{X})}{2\kappa^2} + O(\kappa^{-3} n^{3/2}) \right) \\ &= e^{-t^2/2} + o(1), \end{aligned}$$

since $hn^{3/2}\kappa^{-3} = O(h^{-1/2}) \rightarrow 0$. Here we used the bound on the third moment of $\mathbf{X}$ in a gadget given by Lemma 12. This proves that $\bar{\mathbf{X}} \Rightarrow N(0, 1)$.

Let $\mathcal{L}$ denote the lattice $\langle \mathbf{X} \rangle_{\hat{H}_s^G, Y, \tau} + \mathbb{Z}/\kappa$. Let $\mathcal{N}(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$. We want to show that

$$\sup_{x \in \mathcal{L}} \left| \kappa \mu_{\hat{H}_s^G, Y, \tau}(\bar{\mathbf{X}} = x) - \mathcal{N}(x) \right| = o(1),$$

since $\kappa = \Theta(\sqrt{nh})$. Using Fourier inversion (as in Appendix A), we have for any $K > 0$,

$$\begin{aligned} &2\pi \sup_{x \in \mathcal{L}} \left| \kappa \mu_{G, \beta, \lambda}(\bar{\mathbf{X}} = x) - \mathcal{N}(x) \right| \\ &= \sup_{x \in \mathcal{L}} \left| \int_{-\pi\kappa}^{\pi\kappa} \phi_{\bar{\mathbf{X}}}(t) e^{-itx} dt - \int_{-\infty}^{\infty} e^{-t^2/2 - itx} dt \right| \\ &\leq \int_{-\pi\kappa}^{\pi\kappa} \left| \phi_{\bar{\mathbf{X}}}(t) - e^{-t^2/2} \right| dt + \int_{|t| > \pi\kappa} e^{-t^2/2} dt \\ &\leq \int_{-K}^K \left| \phi_{\bar{\mathbf{X}}}(t) - e^{-t^2/2} \right| dt + \int_{|t| \geq K} e^{-t^2/2} dt + \int_{|t| \geq K} |\phi_{\bar{\mathbf{X}}}(t)| dt \\ &=: A_1 + A_2 + A_3. \end{aligned}$$

Because $\phi_{\overline{X}}(t) = e^{-t^2/2} + o(1)$, applying the bounded convergence theorem gives that $A_1 \rightarrow 0$ as $n \rightarrow \infty$ for any fixed K , so we can choose n large enough to guarantee $A_1 < \varepsilon/3$. We can pick K large enough to ensure $A_2 < \varepsilon/3$. For A_3 , we use the fact that the portion of the characteristic function coming from the isolated vertices has nice behavior. In particular,

$$\begin{aligned} |\phi_{\overline{X}}(t)| &\leq \left(\frac{1 + e^{it/\kappa}}{2} e^{-t/(2\kappa)} \right)^s \\ &\leq e^{-\frac{t^2 s}{4\kappa^2}} = e^{-\Omega(t^2)} \end{aligned}$$

since $s = \Theta(\kappa^2)$. Then by choosing K large enough again we can make $A_3 < \varepsilon/3$ as well. $\square$

With these ingredients we can prove Theorem [10](#).

Proof of Theorem [10](#). Recall that we assume $\Delta \geq 3$, $\beta > \beta_c(\Delta)$, and that $\eta \in [0, \eta_c)$. Let G be the gadget graph that satisfies the conclusions of Lemmas [11](#) and [12](#), and recall the notation of Section [3.1](#) which includes the graph H_s^G on N vertices formed from copies of G and s isolated vertices.

Let b be the value of γ -MEBC on the graph H , and let $\varepsilon > 0$. We will obtain upper and lower bounds on b in terms of $Z_{H_s^G}^{\text{fix}}(\beta, 2\ell - N)$ such that for suitably small ζ , an e^{N^ζ} -relative approximation to $Z_{H_s^G}^{\text{fix}}(\beta, 2\ell - N)$ constrains b to an interval of length at most $h^{2-\varepsilon}$.

Since β and the magnetization $2\ell - N$ are fixed, we will write $Z_{H_s^G}^{\text{fix}}$ for $Z_{H_s^G}^{\text{fix}}(\beta, 2\ell - N)$. Moreover, for a phase vector $Y \in \Sigma_H$, we write $Z_{H_s^G}^{\text{fix}}(Y)$ for the contribution to $Z_{H_s^G}^{\text{fix}}$ from spin assignments with phase vector Y . For $\tau \in \Sigma_R$ we write $Z_{H_s^G}^{\text{fix}}(Y, \tau)$ for the contribution to $Z_{H_s^G}^{\text{fix}}(Y)$ from spin assignments σ which agree with τ on the terminals R . Similarly, since we only consider the usual Ising model with no external field we write $Z_{\hat{H}_s^G}$ for $Z_{\hat{H}_s^G}(\beta, 1)$.

We start by bounding the partition function $Z_{H_s^G}^{\text{fix}}$ from above. The first step is to split the partition function into sums over Y according to whether $M(Y) = M^*$. We have

$$Z_{H_s^G}^{\text{fix}} = \sum_{Y \in \Sigma_H} Z_{H_s^G}^{\text{fix}}(Y) = \sum_{Y: M(Y) = M^*} Z_{H_s^G}^{\text{fix}}(Y) + \sum_{Y: M(Y) \neq M^*} Z_{H_s^G}^{\text{fix}}(Y).$$

For an arbitrary phase vector Y we split $Z_{H_s^G}^{\text{fix}}(Y)$ into a sum over spin assignments $\tau \in \Sigma_R$ to the terminals and pull out the factor of the summand contributed by edges in $\mathcal{E}$, giving

$$Z_{H_s^G}^{\text{fix}}(Y) = \sum_{\tau \in \Sigma_R} Z_{\hat{H}_s^G}^{\text{fix}}(Y, \tau) \prod_{uv \in \mathcal{E}} e^{\frac{\beta}{2} \tau_u \tau_v}.$$

To handle the fixed-magnetization constraint, observe that when σ is drawn from the Ising model $\mu_{\hat{H}_s^G, Y, \tau}$ we have

$$Z_{\hat{H}_s^G}^{\text{fix}}(Y, \tau) = Z_{\hat{H}_s^G}(Y, \tau) \cdot \mu_{\hat{H}_s^G, Y, \tau}(\mathbf{X} = \ell),$$

which we can control with Lemmas [13](#) and [14](#). In the case $M(Y) = M^*$ we have

$$Z_{\hat{H}_s^G}^{\text{fix}}(Y, \tau) = Z_{\hat{H}_s^G}(Y, \tau) \cdot \Omega(1/\sqrt{nh}),$$

and in the case $M(Y) \neq M^*$ we use

$$Z_{\hat{H}_s^G}^{\text{fix}}(Y, \tau) = Z_{\hat{H}_s^G}(Y, \tau) \cdot \exp(-\Omega(n/h)).$$

For the sum over Y with $M(Y) = M^*$ this means for some constant $C > 0$,

$$\sum_{Y: M(Y)=M^*} Z_{\hat{H}_s^G}^{\text{fix}}(Y) \leq \frac{C}{\sqrt{nh}} \sum_{\tau \in \Sigma_R} Z_{\hat{H}_s^G}(Y, \tau) \prod_{uv \in \mathcal{E}} e^{\frac{\beta}{2} \tau_u \tau_v}. \quad (3)$$

Now we can apply the phase-conditioned, nearly-independent terminal spins property of the gadget. Using Lemma [11\(i\)](#) for the inequality (and the fact that $(1 + O(n^{-2\theta}))^h = 1 + o(1)$), we have

$$Z_{\hat{H}_s^G}(Y, \tau) = Z_{\hat{H}_s^G}(Y) \cdot \mu_{\hat{H}^G, Y}(\sigma_R = \tau) \leq (1 + o(1)) Z_{\hat{H}_s^G}(Y) Q_R^Y(\tau),$$

where $Q_R^Y(\tau)$ is the probability measure on Σ_R such that

$$Q_R^Y(\tau) = \prod_{x \in V(H)} Q_{R^x}^{Y_x}(\tau_{R^x}).$$

Continuing from [\(3\)](#) and absorbing factors into the constant, we have

$$\sum_{Y: M(Y)=M^*} Z_{\hat{H}_s^G}^{\text{fix}}(Y) \leq \frac{C}{\sqrt{nh}} \sum_{Y: M(Y)=M^*} Z_{\hat{H}_s^G}(Y) \sum_{\tau \in \Sigma_R} Q_R^Y(\tau) \prod_{uv \in \mathcal{E}} e^{\frac{\beta}{2} \tau_u \tau_v},$$

and the final sum over τ can be expressed in terms of the number $\text{cut}(Y)$ of edges of H which are cut by the phase vector Y . This observation appears in [\[Sly10\]](#) and is precisely why nearly-independent phase-correlated spins are important in reductions such as these.

Recall q defined in [\(2\)](#). For every edge $xy \in E(H)$ cut by Y , there are precisely $2k$ edges in $\mathcal{E}$ such that the measure Q_W^Y gives one endpoint spin $+1$ with probability q and the other endpoint spin $+1$ with probability $1 - q$. Such edges are monochromatic with probability $2q(1 - q)$. Similarly, for every edge $xy \in E(H)$ not cut by Y there are precisely $2k$ edges in $\mathcal{E}$ which are monochromatic with probability $q^2 + (1 - q)^2$. Then for constants $\Theta = 2q(1 - q)e^{\beta/2} + (q^2 + (1 - q)^2)e^{-\beta/2}$ and $\Gamma = 2q(1 - q)e^{-\beta/2} + (q^2 + (1 - q)^2)e^{\beta/2}$ we have

$$\sum_{\tau \in \Sigma_R} Q^Y(\tau) \prod_{uv \in \mathcal{E}} e^{\frac{\beta}{2} \tau_u \tau_v} = \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2k \text{cut}(Y)}.$$

Note that $\Theta < \Gamma$ so that smaller cuts give larger quantities above. Finishing the upper bound started in [\(3\)](#), we have

$$\begin{aligned} \sum_{Y: M(Y)=M^*} Z_{\hat{H}_s^G}^{\text{fix}}(Y) &\leq \frac{C}{\sqrt{nh}} \sum_{Y: M(Y)=M^*} Z_{\hat{H}_s^G}(Y) \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2k \text{cut}(Y)} \\ &\leq \frac{C}{\sqrt{nh}} \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2kb} Z_{\hat{H}_s^G}, \end{aligned}$$

because the γ -MEBC b of H gives the largest contribution $(\Theta/\Gamma)^{2k \text{cut}(Y)}$, and the partition function $Z_{\hat{H}_s^G}$ is an upper bound on $\sum_{Y: M(Y)=M^*} Z_{\hat{H}_s^G}(Y)$.

For phase vectors Y with $M(Y) \neq M^*$ it suffices to consider the worst-case contribution from edges between gadgets. For such Y , $\left| \frac{s}{2} + 2n \sum_{x \in V(H)} \alpha^{Y_x} - \ell \right| > \delta n$, and so we can apply Lemma [13](#) to give

$$\begin{aligned} \sum_{Y: M(Y) \neq M^*} Z_{H_s^G}^{\text{fix}}(Y) &\leq 2e^{-\frac{\delta n}{4h}} \sum_{\tau \in \Sigma_R} Z_{\hat{H}_s^G}(Y, \tau) \prod_{uv \in \mathcal{E}} e^{\frac{\beta}{2} \tau_u \tau_v} \\ &\leq e^{-\Omega(n/h)} e^{\frac{\beta}{2} |\mathcal{E}|} \sum_{\tau \in \Sigma_R} Z_{\hat{H}_s^G}(Y, \tau) \\ &\leq e^{\beta k |E(H)| - \Omega(n/h)} Z_{\hat{H}_s^G}, \end{aligned}$$

because $|\mathcal{E}| = 2k|E(H)|$. Our construction ensures that $k|E(H)| \leq kh^2 = o(n/h)$ so that this is a negligible fraction of $Z_{\hat{H}_s^G}$. Combining these bounds, for all large enough n we have

$$Z_{H_s^G}^{\text{fix}} \leq \left(\frac{C}{\sqrt{nh}} \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2kb} + e^{\beta k |E(H)| - \Omega(n/h)} \right) Z_{\hat{H}_s^G}.$$

The term $\Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2kb}$ is smallest when $b = |E(H)|$, but even in this case it is still at least $e^{-\beta k |E(H)|}$ as $\Theta > e^{-\beta/2}$. Thus, we can absorb the ‘error’ term arising from phase vectors Y with $M(Y) \neq M^*$ into C :

$$Z_{H_s^G}^{\text{fix}} \leq \frac{C}{\sqrt{nh}} \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2kb} Z_{\hat{H}_s^G}. \quad (4)$$

To give a lower bound on $Z_{H_s^G}^{\text{fix}}$ it suffices to consider a single phase vector Y^* with $\text{cut}(Y) = b$ that corresponds to the γ -MEBC of H . Then $M(Y^*) = M^*$ and for some constant $C' > 0$ (which will absorb $(1 + o(1))$ factors in the calculation below), we have

$$\begin{aligned} Z_{H_s^G}^{\text{fix}} &\geq Z_{H_s^G}^{\text{fix}}(Y^*) = \sum_{\tau \in \Sigma_R} Z_{\hat{H}_s^G}^{\text{fix}}(Y^*, \tau) \prod_{uv \in \mathcal{E}} e^{\frac{\beta}{2} \tau_u \tau_v} \\ &\geq \frac{C'}{\sqrt{nh}} \sum_{\tau \in \Sigma_R} Z_{\hat{H}_s^G}(Y^*, \tau) \prod_{uv \in \mathcal{E}} e^{\frac{\beta}{2} \tau_u \tau_v} \\ &\geq \frac{C'}{\sqrt{nh}} Z_{\hat{H}_s^G}(Y^*) \sum_{\tau \in \Sigma_R} Q^{Y^*}(\tau) \prod_{uv \in \mathcal{E}} e^{\frac{\beta}{2} \tau_u \tau_v} \\ &= \frac{C'}{\sqrt{nh}} Z_{\hat{H}_s^G}(Y^*) \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2kb}, \end{aligned}$$

where we apply Lemma [14](#) to obtain the second line and the lower bound in Lemma [12\(i\)](#) to obtain the third. Finally, since we have perfect symmetry between the phases we have $Z_{\hat{H}_s^G}(Y^*) = 2^{-h} Z_{\hat{H}_s^G}$ and

$$Z_{H_s^G}^{\text{fix}} \geq \frac{C' 2^{-h}}{\sqrt{nh}} \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2kb} Z_{\hat{H}_s^G}.$$

The upper bound from [\(4\)](#) and the lower bound above combine to give

$$\frac{C' 2^{-h}}{\sqrt{nh}} \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2kb} \leq \frac{Z_{H_s^G}^{\text{fix}}}{Z_{\hat{H}_s^G}} \leq \frac{C}{\sqrt{nh}} \Gamma^{2k|E(H)|} (\Theta/\Gamma)^{2kb},$$

which provides the bounds

$$T - \frac{\log(2^h/C')}{2k \log(\Gamma/\Theta)} \leq b \leq T + \frac{\log C}{2k \log(\Gamma/\Theta)}$$

on the min-bisection b of H , where

$$T = \frac{\log(Z_{\hat{H}_s^G}/Z_{H_s^G}^{\text{fix}}) + 2k|E(H)| \log \Gamma - \log \sqrt{nh}}{2k \log(\Gamma/\Theta)}.$$

We can approximate $Z_{\hat{H}_s^G}$ to within an absolute constant factor in (randomized) time polynomial in N (which is polynomial in h) by Theorem 9. For the theorem we suppose that we have a relative e^{N^ζ} -approximation of $Z_{H_s^G}^{\text{fix}}$, and hence if $\tilde{T}$ is given by the definition of T above with $Z_{\hat{H}_s^G}$ and $Z_{H_s^G}^{\text{fix}}$ replaced by these approximate values, we have $|\tilde{T} - T| \leq O(N^\zeta/k)$ and hence

$$\tilde{T} - O\left(\frac{N^\zeta + h}{k}\right) \leq b \leq \tilde{T} + O\left(\frac{N^\zeta}{k}\right).$$

Since $k = \Theta(h^3)$, this constrains b to an interval of length $O(N^\zeta h^{-3})$. Using $N = O(nh)$ with $n = O(h^{\theta/4})$, it suffices to choose ζ small enough in terms of θ and ε that $(1 + 4/\theta)\zeta < 5 - \varepsilon$.

Note that we could reduce from solving min-bisection exactly at the cost of weaker inapproximability for Z^{fix} in the proof. If ζ is chosen such that $(1 + 4/\theta)\zeta < 3$ then the bounds constrain the integer b to an interval of length $o(1)$ and hence for large enough h we can find b exactly. $\square$

3.4 Proof of Lemma 12

We prove the lemma in the case of the $+$ phase as the $-$ phase is the same. Note that we can ignore the contribution of vertices in R_0 and the attached trees to $\mathbf{X}$ in the bounds since there are $o(n^{1/2})$ of these vertices. So for this section $\mathbf{X}$ and $X(\sigma)$ will refer to the number of $+$ spins in the vertices of G' .

We first prove the three bounds of the lemma without conditioning on $\{\sigma_{R_0} = \tau\}$. We will show that for large enough n , with probability at least $8/10$ over the choice of gadget, we have the following bounds:

$$\langle |\mathbf{X} - 2n\alpha^+| \rangle_{G,+} = O(\sqrt{n}) \quad (5)$$

$$\langle |\mathbf{X} - 2n\alpha^+|^2 \rangle_{G,+} = O(n) \quad (6)$$

$$\langle |\mathbf{X} - 2n\alpha^+|^3 \rangle_{G,+} = O(n^{3/2}) \quad (7)$$

$$\langle e^{t_0(\mathbf{X} - 2\alpha^+n)} \rangle_{G,+} \leq e^{O(t_0^2 n)}. \quad (8)$$

Let $\xi : \mathbb{R} \rightarrow \mathbb{R}$ be a non-negative function that satisfies $\xi(x) \leq e^{c|x|}$ for some constant $c > 0$. We aim to prove bounds on $\langle \xi(\mathbf{X} - 2\alpha^+n) \rangle_{G,+}$ for four choices of functions ξ : $\xi(x) = |x|^k$ for $k \in \{1, 2, 3\}$, and $\xi(x) = e^{t_0 x}$.

For such a function ξ , we can write

$$\langle \xi(\mathbf{X} - 2\alpha^+ n) \rangle_{G,+} \leq \frac{\sum_{\alpha \geq 1/2} \xi(2\alpha n - 2\alpha^+ n) Z_{G,\alpha}}{Z_{G,+}}$$

(this is an inequality instead of an equality simply because we include all configurations with $\alpha = 1/2$).

By (iii) of Lemma 11 we have $Z_{G,+} > \frac{1}{C} \mathbb{E} Z_{G,+}$ with probability at least $1 - 1/10$. By Markov's inequality we have

$$\sum_{\alpha \geq 1/2} \xi(2\alpha n - 2\alpha^+ n) Z_{G,\alpha} \leq 100 \sum_{\alpha \geq 1/2} \xi(2\alpha n - 2\alpha^+ n) \mathbb{E} Z_{G,\alpha}$$

for all four choices of ξ with probability at least $1 - 4/100$. Thus with probability at least $1 - 1/10 - 4/100 \geq 8/10$ we have

$$\langle \xi(\mathbf{X} - 2\alpha^+ n) \rangle_{G,+} \leq 100C \frac{\sum_{\alpha \geq 1/2} \xi(2\alpha n - 2\alpha^+ n) \mathbb{E} Z_{G,\alpha}}{\mathbb{E} Z_{G,+}},$$

so to prove (5), (6), (7), (8) it is enough to show that $\frac{\sum_{\alpha \geq 1/2} \xi(2\alpha n - 2\alpha^+ n) \mathbb{E} Z_{G,\alpha}}{\mathbb{E} Z_{G,+}}$ satisfies the desired bounds.

Now using the bound

$$\frac{\mathbb{E} Z_{G,\alpha}}{\mathbb{E} Z_{G,+}} = O(n^{-1/2}) e^{-\Omega(n(\alpha - \alpha^+)^2)}$$

from Lemma 11, we can bound

$$\begin{aligned} \frac{\sum_{\alpha \geq 1/2} \xi(2\alpha n - 2\alpha^+ n) \mathbb{E} Z_{G,\alpha}}{\mathbb{E} Z_{G,+}} &\leq O(n^{-1/2}) \sum_{n(1-2\alpha^+) \leq \ell \leq 2n(1-\alpha^+)} \xi(\ell) e^{-\Omega(\ell^2/n)} + o(1) \\ &= O(n^{-1/2}) \int_{n(1-2\alpha^+)}^{2n(1-\alpha^+)} \xi(u) e^{-\Omega(u^2/n)} du \\ &= O(n^{1/2}) \int_{1-2\alpha^+}^{2(1-\alpha^+)} \xi(un) e^{-\Omega(nx^2)} dx \\ &= O\left(n^{1/2} \int_{-\infty}^{\infty} \xi(xn) e^{-\Omega(nx^2)} dx\right). \end{aligned}$$

Since

$$\int_{-\infty}^{\infty} |xn|^k e^{-\Omega(nx^2)} dx = O(n^{(k-1)/2})$$

and

$$\int_{-\infty}^{\infty} e^{t_0 nx} e^{-\Omega(nx^2)} dx = e^{O(t_0^2 n)}$$

we obtain (5), (6), (7), (8).

Now we transfer these bounds to the measure conditioned on $\{\sigma_{R_0} = \tau\}$. For the moment generating function we have

$$\begin{aligned} \langle e^{t_0(\mathbf{X} - 2\alpha^+ n)} \rangle_{G,+, \tau} &= \frac{\langle e^{t_0(\mathbf{X} - 2\alpha^+ n)} \mathbf{1}_{\sigma_{R_0} = \tau} \rangle_{G,+}}{\mu_{G,+}(\sigma_{R_0} = \tau)} \\ &\leq \frac{\langle e^{t_0(\mathbf{X} - 2\alpha^+ n)} \rangle_{G,+}}{Q_{R_0}^+(\tau)(1 - O(n^{-2\theta}))} \\ &\leq e^{O(|R_0|)} e^{O(t_0^2 n)} \\ &= e^{O(t_0^2 n)}, \end{aligned}$$

where we used that $t_0^2 n = O(nh^{-2}) = O(n^{1-\theta/2})$ and $|R_0| = O(n^\theta) = o(n^{1-\theta/2})$.

For the k th moment,

$$\begin{aligned} \langle |\mathbf{X} - 2n\alpha^+|^k \rangle_{G,+, \tau} &= \frac{\langle |\mathbf{X} - 2n\alpha^+|^k \cdot \mathbf{1}_{\sigma_{R_0} = \tau} \rangle_{G,+}}{\mu_{G,+}(\sigma_{R_0} = \tau)} \\ &\leq \frac{\langle |\mathbf{X} - 2n\alpha^+|^k \cdot \mathbf{1}_{\sigma_{R_0} = \tau} \cdot \mathbf{1}_{\sigma_{W_0} \in \mathcal{B}^c} \rangle_{G,+} + \langle |\mathbf{X} - 2n\alpha^+|^k \cdot \mathbf{1}_{\sigma_{W_0} \in \mathcal{B}} \rangle_{G,+}}{Q_{R_0}^+(\tau)(1 - O(n^{-2\theta}))} \\ &\leq \frac{\langle |\mathbf{X} - 2n\alpha^+|^k \cdot \mathbf{1}_{\sigma_{R_0} = \tau} \cdot \mathbf{1}_{\sigma_{W_0} \in \mathcal{B}^c} \rangle_{G,+} + (2n)^k \exp(-n^{2\theta})}{Q_{R_0}^+(\tau)(1 - O(n^{-2\theta}))} \\ &= \frac{\langle |\mathbf{X} - 2n\alpha^+|^k \cdot \mathbf{1}_{\sigma_{R_0} = \tau} \cdot \mathbf{1}_{\sigma_{W_0} \in \mathcal{B}^c} \rangle_{G,+}}{Q_{R_0}^+(\tau)(1 - O(n^{-2\theta}))} + o(1) \\ &= \frac{\sum_{\sigma \in \Sigma_G} |X(\sigma) - 2n\alpha^+|^k \mathbf{1}_{\sigma_{R_0} = \tau} \cdot \mathbf{1}_{\sigma_{W_0} \in \mathcal{B}^c} \cdot \mu_{G,+}(\sigma)}{Q_{R_0}^+(\tau)(1 - O(n^{-2\theta}))} + o(1) \\ &= \frac{\sum_{\sigma \in \Sigma_G} |X(\sigma) - 2n\alpha^+|^k \mu(\sigma_{R_0} = \tau | \sigma_{W_0}) \cdot \mathbf{1}_{\sigma_{W_0} \in \mathcal{B}^c} \cdot \mu_{G,+}(\sigma)}{Q_{R_0}^+(\tau)(1 - O(n^{-2\theta}))} + o(1) \\ &= \frac{Q_{R_0}^+(\tau)(1 + O(n^{-3\theta})) \langle |\mathbf{X} - 2n\alpha^+|^k \cdot \mathbf{1}_{\sigma_{W_0} \in \mathcal{B}^c} \rangle_{G,+}}{Q_{R_0}^+(\tau)(1 - O(n^{-2\theta}))} + o(1) \\ &\leq (1 + O(n^{-2\theta})) \langle |\mathbf{X} - 2n\alpha^+|^k \rangle_{G,+} + o(1) \\ &= O(n^{(k-1)/2}), \end{aligned}$$

where we have used from Lemma 11 that for all $\tau_{W_0} \in \mathcal{B}^c$,

$$\mu_{G,+}(\sigma_{R_0} = \tau | \sigma_{W_0} = \tau_{W_0}) = Q_{R_0}^+(\tau)(1 + O(n^{-3\theta})),$$

and we have used the fact that conditioned on σ_{W_0} , $X(\sigma)$ and σ_{R_0} are independent (using here that in this section $X(\sigma)$ only counts + spins to the vertices of G').

4 Extremal bounds on the mean magnetization

The proof of Theorem 3 is based on that of Krinsky in [Kri75] which applies to lattices such as $\mathbb{Z}^d$. We require some simple calculus facts recorded in the following lemma.

Lemma 15. *Let $\beta > 0$, $h_1(x) = \operatorname{artanh}(\tanh x \tanh \frac{\beta}{2})$, and let $h_2(x, y) = \tanh(x + \operatorname{artanh}(\tanh y \tanh \frac{\beta}{2}))$. Then h_1 is strictly concave on $(0, \infty)$, h_2 is strictly concave when $x, y \in (0, \infty)$, and $h_2(x, x)$ is an increasing function of x .*

Proof. For the first statement, note that

$$h_1''(x) = -\frac{2 \sinh \beta \sinh(2x)}{(\cosh \beta + \cosh(2x))^2}.$$

Now let $g = \operatorname{artanh}$. For the second statement, note that the Hessian matrix of h_2 has determinant

$$\frac{\sinh \beta \sinh(2y) \tanh(x + g(\tanh y \tanh \frac{\beta}{2}))}{\cosh^2(\frac{\beta}{2} - y) \cosh^2(\frac{\beta}{2} + y) \cosh^4(x + g(\tanh y \tanh \frac{\beta}{2}))} > 0,$$

and

$$\frac{\partial^2}{\partial x^2} h_2(x, y) = -\frac{2 \sinh(g(\tanh y \tanh \frac{\beta}{2}) + x)}{\cosh^3(g(\tanh y \tanh \frac{\beta}{2}) + x)} < 0.$$

This means that the Hessian is negative definite when $x, y \in (0, \infty)$. Finally, observe that

$$\frac{\partial}{\partial x} h_2(x, x) = \left(\frac{\sinh \beta}{\cosh \beta + \cosh(2x)} + 1 \right) \operatorname{sech}^2(x + g(\tanh x \tanh \frac{\beta}{2})) > 0. \quad \square$$

Proof of Theorem 3. Let $G = (V, E)$ be a graph and consider the ferromagnetic Ising model with partition function

$$Z_G(\vec{\beta}, \lambda) = \sum_{\sigma \in \Sigma_G} e^{\frac{1}{2} \sum_{uv \in E(G)} \beta_{uv} \sigma_u \sigma_v} \lambda^{M(\sigma)},$$

where we allow each edge $uv \in E$ to have its own inverse temperature parameter β_{uv} . Specializing to $\beta_{uv} = \beta$ for all edges uv , we recover the definition used elsewhere in this work.

When β and λ are understood from context, given a function f with domain Σ_G , let $\langle f \rangle_G$ be the expected value of f with respect to the Ising model on G . We also write $\langle f \rangle_{G-uv}$ for the expected value of f with respect to the Ising model on the graph formed from G by removing the edge uv , which is equivalent to setting the parameter β_{uv} to zero. We extend this notation to $\langle f \rangle_{G-F}$ when we want to remove some set F of edges.

A key feature of the ferromagnetic Ising model with non-negative external field is the following list of Griffiths' inequalities [Gri67], also known as the GKS inequalities after Kelly and Sherman who generalized Griffiths' work [KS68]. For $A \subset V$, let $\sigma_A = \prod_{v \in A} \sigma_v$. Then we have for any graph $G = (V, E)$, $A, B \subset V$, and $uv \in E$,

$$\langle \sigma_A \rangle \geq 0 \tag{9}$$

$$\langle \sigma_A \sigma_B \rangle - \langle \sigma_A \rangle \langle \sigma_B \rangle \geq 0 \tag{10}$$

$$\langle \sigma_A \rangle - \langle \sigma_A \rangle_{G-uv} \geq 0. \tag{11}$$

In fact, (I0) implies (I1) by considering the derivative of $\langle \sigma_A \rangle$ with respect to β_{uv} . With $B = \{u, v\}$ we have

$$\frac{\partial}{\partial \beta_{uv}} \langle \sigma_A \rangle = \frac{1}{2} \langle \sigma_A \sigma_B \rangle - \frac{1}{2} \langle \sigma_A \rangle \langle \sigma_B \rangle \geq 0,$$

where the inequality is by (I0).

We will apply Griffiths' inequalities during some careful manipulation of expectations using the following identities. For $s = \pm 1$ and any β we have

$$e^{\frac{\beta}{2}s} = \left(1 + s \tanh \frac{\beta}{2}\right) \cosh \frac{\beta}{2}, \quad (12)$$

which follows from the definitions of the hyperbolic functions in terms of exponential functions. We also use the addition formula

$$\operatorname{artanh} \left(\frac{x+y}{1+xy} \right) = \operatorname{artanh} x + \operatorname{artanh} y.$$

From now on, we work with $\beta_{uv} = \beta$ for all $uv \in E$. Let $u \in V$ and $v, w \in N(u)$ with $v \neq w$. Applying (I2) to the term $e^{\frac{\beta}{2}\sigma_u \sigma_v}$ which occurs in both the numerator and the denominator of $\langle \sigma_u \rangle$ gives

$$\langle \sigma_u \rangle = \frac{\langle \sigma_u \rangle_{G-uv} + \langle \sigma_v \rangle_{G-uv} \tanh \frac{\beta}{2}}{1 + \langle \sigma_u \sigma_v \rangle_{G-uv} \tanh \frac{\beta}{2}}, \quad (13)$$

and the same identity applied to the edge uw in $G - uv$ gives

$$\langle \sigma_u \rangle_{G-uv} = \frac{\langle \sigma_u \rangle_{G-uv-uw} + \langle \sigma_w \rangle_{G-uv-uw} \tanh \frac{\beta}{2}}{1 + \langle \sigma_u \sigma_w \rangle_{G-uv-uw} \tanh \frac{\beta}{2}}. \quad (14)$$

To each of these we apply (I0) to the expectation in the denominator, giving

$$\langle \sigma_u \rangle \leq \frac{\langle \sigma_u \rangle_{G-uv} + \langle \sigma_v \rangle_{G-uv} \tanh \frac{\beta}{2}}{1 + \langle \sigma_u \rangle_{G-uv} \langle \sigma_v \rangle_{G-uv} \tanh \frac{\beta}{2}}, \text{ and} \quad (15)$$

$$\langle \sigma_u \rangle_{G-uv} \leq \frac{\langle \sigma_u \rangle_{G-uv-uw} + \langle \sigma_w \rangle_{G-uv-uw} \tanh \frac{\beta}{2}}{1 + \langle \sigma_u \rangle_{G-uv-uw} \langle \sigma_w \rangle_{G-uv-uw} \tanh \frac{\beta}{2}}. \quad (16)$$

Applying $g = \operatorname{artanh}$ to both sides of (I5) and using the addition formula, we obtain

$$g(\langle \sigma_u \rangle) \leq g(\langle \sigma_u \rangle_{G-uv}) + g(\langle \sigma_v \rangle_{G-uv} \tanh \frac{\beta}{2}). \quad (17)$$

Doing this again with (I6), we also use (I1) with the edge uw in the graph $G - uv$, giving $\langle \sigma_w \rangle_{G-uv-uw} \leq \langle \sigma_w \rangle_{G-uw}$ and hence

$$g(\langle \sigma_u \rangle_{G-uv}) \leq g(\langle \sigma_u \rangle_{G-uv-uw}) + g(\langle \sigma_w \rangle_{G-uv-uw} \tanh \frac{\beta}{2}). \quad (18)$$

Observe that we can iterate the process used to obtain (I8) over each $w \in N(u) \setminus \{v\}$ in turn to obtain

$$g(\langle \sigma_u \rangle_{G-uv}) \leq \log \lambda + \sum_{w \in N(u) \setminus \{v\}} g(\langle \sigma_w \rangle_{G-uw} \tanh \frac{\beta}{2}), \quad (19)$$

where we have used the fact that removing the set $F_u := \{uv : v \in N(u)\}$ of edges incident to u we have

$$\langle \sigma_u \rangle_{G-F_u} = \frac{\lambda - \lambda^{-1}}{\lambda + \lambda^{-1}} = \tanh(\log \lambda)$$

because u is an isolated vertex in $G - F_u$.

At this point, Krinsky assumes that G is both edge and vertex transitive so that for some $L > 0$ and for all $uv \in E$ we have

$$\langle \sigma_u \rangle_{G-uv} = \langle \sigma_v \rangle_{G-uv} = \tanh L.$$

In this special case, (19) becomes

$$L \leq \log \lambda + (\Delta - 1)g(\tanh L \tanh \frac{\beta}{2}),$$

where Δ is the degree of a vertex in G . This implies that L is bounded above by the largest solution L^* to

$$L^* = \log \lambda + (\Delta - 1)g(\tanh L^* \tanh \frac{\beta}{2}). \quad (20)$$

Plugging this into (17) and observing that for any vertex-transitive graph $\langle \sigma_u \rangle$ is the mean magnetization η_G , we have

$$\eta_G \leq \tanh(L^* + g(\tanh L^* \tanh \frac{\beta}{2})).$$

The right-hand side is precisely $\eta_{\Delta, \beta, \lambda}^+$, the mean magnetization of the $+$ measure on the infinite Δ -regular tree, which one can derive from first principles (as in, e.g., [Bax82]). In fact, it suffices to observe that every inequality we applied to obtain this bound holds with equality in the tree. That is, in the infinite Δ -regular tree $\langle \sigma_u \sigma_v \rangle_{G-uv} = \langle \sigma_u \rangle_{G-uv} \langle \sigma_v \rangle_{G-uv}$ since removing uv leaves u and v in different connected components so σ_u and σ_v are independent. Similarly, we have $\langle \sigma_u \sigma_w \rangle_{G-uv-uw} = \langle \sigma_u \rangle_{G-uv-uw} \langle \sigma_w \rangle_{G-uv-uw}$ and since w is in a different component from the edge uv after uw is removed, $\langle \sigma_w \rangle_{G-uv-uw} = \langle \sigma_w \rangle_{G-uw}$. The tree is edge and vertex transitive, proving that the derived upper bound is given by the mean magnetization of some measure on the tree. As the $+$ measure stochastically dominates all other translation-invariant measures on the tree, it corresponds to the largest solution L^* to (20).

We now will apply this argument to a finite graph G that is not necessarily vertex or edge transitive. First, we observe that we can reduce to the case that G is regular by a well-known construction. Suppose that G has maximum degree Δ but minimum degree $\delta \leq \Delta - 1$. We construct a graph H with minimum degree $\delta + 1$ such that $\eta_G \leq \eta_H$. Let H_0 be formed from the disjoint union of two copies of G , so that the mean magnetization of H_0 is equal to that of G . For $i \geq 0$, if there is a vertex u of degree δ in H_i let H_{i+1} be formed from H_i by connecting u to its copy in the other copy of G . The inequality (11) shows that η_{H_i} is non-decreasing as i increases because adding the edge can only increase any term $\langle \sigma_v \rangle$, and the mean magnetization is the average of these terms over all vertices v . When the process terminates at some H , the minimum degree of H is $\delta + 1$, and we cannot have decreased the mean magnetization. Iterating this construction, we can obtain a Δ -regular graph H whose mean magnetization is an upper bound on the mean magnetization of G , hence it suffices to prove the theorem in the case of a Δ -regular graph.

In a Δ -regular graph, careful averaging and applications of Jensen's inequality allow us to recover the result obtained for edge and vertex transitive graphs. We can interpret (19) as a property of an oriented edge $\vec{uv}$, and average over all edges incident to u oriented away from u . Let $L_{\vec{uv}}$ be given by $\tanh L_{\vec{uv}} = \langle \sigma_u \rangle_{G-uv}$, so that averaging (19) over $v \in N(u)$ gives

$$\frac{1}{\Delta} \sum_{v \in N(u)} L_{\vec{uv}} \leq \log \lambda + \frac{\Delta-1}{\Delta} \sum_{v \in N(u)} g(\tanh L_{\vec{uv}} \tanh \frac{\beta}{2}). \quad (21)$$

By Lemma 15, the function $x \mapsto g(\tanh x \tanh \frac{\beta}{2})$ is concave on $(0, \infty)$. This means that (21) and Jensen's inequality give

$$\frac{1}{\Delta} \sum_{v \in N(u)} L_{\vec{uv}} \leq \log \lambda + (\Delta-1)g\left(\tanh \left[\frac{1}{\Delta} \sum_{v \in N(u)} L_{\vec{uv}}\right] \tanh \frac{\beta}{2}\right). \quad (22)$$

To clean this up, we define

$$A_u := \frac{1}{\Delta} \sum_{v \in N(u)} L_{\vec{uv}}, \quad B_u := \frac{1}{\Delta} \sum_{v \in N(u)} L_{\vec{vu}},$$

so that (22) gives

$$A_u \leq \log \lambda + (\Delta-1)g\left(\tanh B_u \tanh \frac{\beta}{2}\right). \quad (23)$$

This we average over a uniform random $u \in V$ and again appeal to concavity. Here we finally obtain the desired equation because the averages satisfy

$$\bar{L} := \frac{1}{\Delta n} \sum_{uv \in E} (L_{\vec{uv}} + L_{\vec{vu}}) = \frac{1}{\Delta n} \sum_{u \in V} \sum_{v \in N(u)} L_{\vec{uv}} = \frac{1}{\Delta n} \sum_{u \in V} \sum_{v \in N(u)} L_{\vec{vu}},$$

so an application of Jensen's inequality gives for $\bar{L}$ what we had for L in the case of a transitive graph,

$$\bar{L} \leq \log \lambda + (\Delta-1)g(\tanh \bar{L} \tanh \frac{\beta}{2}).$$

As before, this means that $\bar{L} \leq L^*$.

To conclude the argument in the Δ -regular case, we apply the same averaging trick to (17). For any edge $uv \in E$,

$$g(\langle \sigma_u \rangle) \leq g(\langle \sigma_u \rangle_{G-uv}) + g(\langle \sigma_v \rangle_{G-uv} \tanh \frac{\beta}{2}),$$

so fixing u and averaging over $v \in N(u)$ gives

$$g(\langle \sigma_u \rangle) \leq \frac{1}{\Delta} \sum_{v \in N(u)} g(\langle \sigma_u \rangle_{G-uv}) + \frac{1}{\Delta} \sum_{v \in N(u)} g(\langle \sigma_v \rangle_{G-uv} \tanh \frac{\beta}{2}).$$

Applying Jensen's inequality again, we have

$$g(\langle \sigma_u \rangle) \leq A_u + g(\tanh B_u \tanh \frac{\beta}{2}),$$

and so

$$\langle \sigma_u \rangle \leq \tanh(A_u + g(\tanh B_u \tanh \frac{\beta}{2})). \quad (24)$$

By Lemma 15, the right-hand side is concave as a function of A_u and B_u . Averaging (24) over $u \in V$ and applying Jensen's inequality, we conclude

$$\begin{aligned}\eta_G &= \frac{1}{n} \sum_{u \in V} \langle \sigma_u \rangle \leq \tanh(\bar{L} + g(\tanh \bar{L} \tanh \frac{\beta}{2})) \\ &\leq \tanh(L^* + g(\tanh L^* \tanh \frac{\beta}{2})) = \eta_{\Delta, \beta, \lambda}^+, \end{aligned}$$

using the fact that $x \mapsto \tanh(x + g(\tanh x \tanh \frac{\beta}{2}))$ is non-decreasing proved in Lemma 15. $\square$

5 Algorithms

In this section we prove Theorems 1 and 2(a). By symmetry, it suffices to consider the case when $\eta \geq 0$. We can exclude the trivial case $\eta = 1$ since there is just a single spin configuration in that case.

We will use several ingredients from Section 2. Fix β, η, Δ satisfying the conditions of either theorem, and let G be a graph of maximum degree Δ on n vertices. Let $\ell = \lfloor \frac{1+\eta}{2} n \rfloor$ so that our goal is to sample an Ising configuration σ with $X(\sigma) = \ell$.

Since we can efficiently sample from $\mu_{G, \beta, \lambda}$ for any λ via Theorem 9, we can perform a binary search on values of λ , estimating $\langle \mathbf{X} \rangle_{G, \beta, \lambda}$, to find a λ so that

$$|\langle \mathbf{X} \rangle_{G, \beta, \lambda} - \ell| = o(\sqrt{n}). \quad (25)$$

Given such an activity λ , we will approximately sample from $\mu_{G, \beta, \lambda}$ until we sample a configuration σ with $X(\sigma) = \ell$ and then output σ . For this algorithm to be efficient we must ensure that the probability of hitting this value is not too small; in fact, we will show that it is $\Theta(n^{-1/2})$.

For Theorem 1 this follows immediately from Proposition 7 which provides a local central limit theorem and $\Theta(n)$ variance for $\mathbf{X}$ for $\beta < \beta_c(\Delta)$ and any activity λ .

For Theorem 2, we need to ensure that we can find λ satisfying (25) that is bounded away from 1 independently of n so we can apply Proposition 7. This is guaranteed by the extremal result, Theorem 3, and the conditions of the Theorem 2. In particular, because $\eta > \eta_+(\beta, \Delta)$ (and by continuity of the magnetization of the $+$ measure on the tree) there is some $\lambda_{\min} > 1$ so that $\eta = \eta_{\Delta, \beta, \lambda_{\min}}^+$. Theorem 3 then says that to achieve mean magnetization η on any $G \in \mathcal{G}_\Delta$ we must take $\lambda \geq \lambda_{\min}$, thus giving the required uniform bound away from 1.

In what follows we give the details of the approach. We note that the running time of our algorithm could certainly be improved by using a faster Ising sampler (e.g. [MS13]) or by using the techniques of [JPSS21], but here we will not try to optimize the running time beyond finding polynomial-time algorithms.

The existence of the efficient sampling schemes in Theorems 1 and 2(a) are proved in Theorem 18 in Section 5.2. The existence of approximate counting algorithms follows from the sampling algorithms via a standard reduction. We provide the details in Appendix B.

5.1 Bounds on the activity

Here we prove a lemma guaranteeing the existence of a good activity λ for use in our sampling algorithms. We write $\lambda_{\Delta,\beta}^{-1}(\eta)$ for the value of λ such that $\eta_{\beta,\Delta,\lambda}^+ = \eta$. In particular, when $\eta > \eta_+(\beta, \Delta)$ we have $\lambda_{\Delta,\beta}^{-1}(\eta) > 1$.

Lemma 16. *Let $\Delta \geq 3$, $\beta < \beta_c(\Delta)$, and $\eta \in [0, 1)$ be fixed. Let $\lambda_{\min} = 1$ and $\lambda_{\max} = \sqrt{\frac{1+\eta}{1-\eta}}e^{\beta\Delta/2}$. Then for any $G \in \mathcal{G}_\Delta$ on n vertices, there exists an integer $t \in \{0, \dots, \lfloor (\lambda_{\max} - \lambda_{\min})n \rfloor\}$ so that for $\ell = \lfloor \frac{\eta+1}{2}n \rfloor$ and $\lambda_t = \lambda_{\min} + \frac{t}{n}$ we have*

$$|\langle \mathbf{X} \rangle_{G,\beta,\lambda_t} - \ell| = O(1), \quad (26)$$

where the implied constant depends only on Δ, β, η . The same holds for $\beta > \beta_c$ and $\eta \in (\eta_+(\beta, \Delta), 1)$ with $\lambda_{\min} = \lambda_{\Delta,\beta}^{-1}(\eta) > 1$ and $\lambda_{\max} = \sqrt{\frac{1+\eta}{1-\eta}}e^{\beta\Delta/2}$.

Before we prove Lemma 16, we need one simple bound on the magnetization of a bounded-degree graph.

Lemma 17. *For all $\Delta \geq 1$, $\beta \geq 0$, $\eta \in [0, 1)$, and any $G \in \mathcal{G}_\Delta$, the value of λ such that the mean magnetization $\eta_G(\beta, \lambda)$ is exactly η satisfies*

$$1 \leq \lambda \leq \sqrt{\frac{1+\eta}{1-\eta}}e^{\beta\Delta/2}.$$

Proof. The lower bound follows from symmetry of the Ising model. For the upper bound, let $\sigma \in \Sigma_G$ be drawn from $\mu_{G,\beta,\lambda}$ and suppose that $\eta = \eta_G(\beta, \lambda)$. Let u be a uniform random vertex of G , and let $Y = M(\sigma_{N(u)})$ be the magnetization of the neighbors of u . That is, when u has exactly j neighbors with spin $+$, $Y = 2j - \deg(u)$. Then a direct computation in the Ising model gives

$$\eta = \left\langle \frac{\lambda e^{\frac{\beta}{2}Y} - \lambda^{-1}e^{-\frac{\beta}{2}Y}}{\lambda e^{\frac{\beta}{2}Y} + \lambda^{-1}e^{-\frac{\beta}{2}Y}} \right\rangle_{G,\beta,\lambda} = \left\langle \frac{\lambda^2 - e^{-\beta Y}}{\lambda^2 + e^{-\beta Y}} \right\rangle_{G,\beta,\lambda}.$$

Since $Y \geq -\Delta$ and the function of Y inside the expectation is increasing for $\beta > 0$, we immediately obtain

$$\eta \geq \frac{\lambda^2 - e^{\beta\Delta}}{\lambda^2 + e^{\beta\Delta}} \iff \lambda \leq \sqrt{\frac{1+\eta}{1-\eta}}e^{\beta\Delta/2}. \quad \square$$

We now prove Lemma 16.

Proof of Lemma 16. By Theorem 3 and Lemma 17, the value of λ such that $\eta_G(\beta, \lambda) = \eta$ satisfies $\lambda_{\min} \leq \lambda \leq \lambda_{\max}$. A standard calculation gives that

$$\frac{\partial}{\partial \lambda} \eta_G(\beta, \lambda) = \frac{2}{n\lambda} \text{var}(\mathbf{X}).$$

Then by Proposition 7 we have $\frac{\partial}{\partial \lambda} \eta_G(\beta, \lambda) > 0$ and $\frac{\partial}{\partial \lambda} \eta_G(\beta, \lambda) = O(1)$. This means that $\eta_G(\beta, \lambda)$ and hence $\langle \mathbf{X} \rangle_{G,\beta,\lambda}$ are strictly increasing as functions of λ , and that $\langle \mathbf{X} \rangle_{G,\beta,\lambda}$ can increase by at most $O(1)$ on any interval of length $1/n$. $\square$

5.2 The sampling algorithm

Our algorithm is the combination of a simple binary search on values of λ and repeated sampling from distributions $\hat{\mu}_{\beta,\lambda}$ that approximate the usual Ising model $\mu_{G,\beta,\lambda}$.

Algorithm: Sample- k

- INPUT: $\Delta, \beta, \eta, \varepsilon$ and $G \in \mathcal{G}_\Delta$ on n vertices.
 - OUTPUT: $\sigma \in \Sigma_G(k)$ distributed within ε total variation distance of the fixed-magnetization Ising model $\nu_{G,\beta,k}$, where k is the largest integer such that $k \equiv n \pmod 2$ and $k \leq \eta n$.
1. Let $\lambda_{\min}, \lambda_{\max}$ be as given in Lemma 16.
 2. For $t = 0, \dots, \lfloor (\lambda_{\max} - \lambda_{\min})n \rfloor$, let $\lambda_t = \lambda_{\min} + t/n$.
 3. Let $\Lambda_0 = \{\lambda_t : t = 0, \dots, \lfloor (\lambda_{\max} - \lambda_{\min})n \rfloor\}$.
 4. FOR $i = 1, \dots, C \log n$,
 - (a) Let λ be a median of the set Λ_{i-1} .
 - (b) With $N = C'n^2 \log(\frac{\log n}{\varepsilon})$, take N independent samples $\sigma_1, \dots, \sigma_N$ from a distribution $\hat{\mu}_{G,\beta,\lambda}$ on Σ_G such that $\|\hat{\mu}_{G,\beta,\lambda} - \mu_{G,\beta,\lambda}\|_{TV} < \varepsilon' = \frac{1}{CN \log n}$.
 - (c) If there exists $j \in \{1, \dots, N\}$ so that $M(\sigma_j) = k$, then output σ_j for the smallest such j and HALT.
 - (d) Let $\bar{k} = \frac{1}{N} \sum_{j=1}^N M(\sigma_j)$.
 - (e) If $\bar{k} \leq k$, let $\Lambda_i = \{\lambda' \in \Lambda_{i-1} : \lambda' > \lambda\}$. If instead $\bar{k} > k$, let $\Lambda_i = \{\lambda' \in \Lambda_{i-1} : \lambda' < \lambda\}$.
 5. If no spin assignment of magnetization k has been obtained by the end of the FOR loop (or if $\Lambda_j = \emptyset$ at any step), output a spin assignment of magnetization k by taking the first $(k+n)/2$ vertices in an arbitrary ordering on $V(G)$ and setting their spins to $+$ and remaining spins to $-$.

Theorem 18. *Let $\hat{\nu}_{G,\beta,k}$ be the output distribution of the algorithm Sample- k . Then for n large enough,*

1. $\|\hat{\nu}_{G,\beta,k} - \nu_{G,\beta,k}\|_{TV} < \varepsilon$.
2. *The running time of Sample- k is polynomial in n and $\log(1/\varepsilon)$.*

This proves the approximate sampling portions of Theorems 11 and 2(a).

Proof of Theorem 18. We first give the proof under the assumption that each $\hat{\mu}_{\beta,\lambda}$ is precisely the Ising model $\mu_{G,\beta,\lambda}$, and subsequently we will reduce the general case that $\hat{\mu}_{\beta,\lambda}$ is close to $\mu_{G,\beta,\lambda}$ to this case with a standard coupling argument.

We say a *failure* occurs at step i in the FOR loop if either

- (1) $|n\eta_G(\beta, \lambda) - \bar{k}| > 1/4$, or
- (2) $|n\eta_G(\beta, \lambda) - k| \leq 1/2$ but none of the samples $\sigma_1, \dots, \sigma_N$ have magnetization exactly k and so the algorithm did not HALT on line (d).

Note that avoiding [\(1\)](#) means that the sample mean $\bar{k}$ of the N spin assignments sampled is close to the true mean $n\eta_G(\beta, \lambda)$, and avoiding [\(2\)](#) means that in the case that the true mean magnetization is close to k , we successfully sample a spin assignment of the desired magnetization k . To establish that the probability of a failure occurring at any step is at most $\varepsilon/2$, it suffices to establish that the probability of each type of failure in a given step is at most $\varepsilon/(4C \log n)$. Consider an arbitrary step i with the value of λ assigned for that step in line (a), and note that $\bar{k}$ is the mean of N independent samples from $\hat{\mu}_{\beta, \lambda}$. Under the assumption that $\hat{\mu}_{\beta, \lambda} = \mu_{G, \beta, \lambda}$, we have $\mathbb{E}\bar{k} = n\eta_G(\beta, \lambda)$, so by Hoeffding's inequality we have

$$\mathbb{P}(|n\eta_G(\beta, \lambda) - \bar{k}| > 1/4) \leq 2e^{-N/(32n^2)}.$$

This is at most the desired $\varepsilon/(4C \log n)$ when $N \geq \Omega(n^2 \log(\log(n)/\varepsilon))$.

For the second type of failure, we suppose that the current value of λ means that $|n\eta_G(\beta, \lambda) - k| \leq O(1)$, but that none of the N samples from $\hat{\mu}_{\beta, \lambda}$ give a state with magnetization exactly k . Each 'trial' to get a state of magnetization k is independent and succeeds with probability $p \geq \Omega(1/\sqrt{n})$ by Lemma [7](#). Then we have no successful trials with probability $(1 - p)^N$, which is at most $\varepsilon/(4C \log n)$ when $N \geq \Omega(\sqrt{n} \log(\log(n)/\varepsilon))$.

The above lower bounds on N show that the value given in line (b) suffices. With a bound on $\varepsilon/2$ on the probability of any failure, we now show that the output state has distribution within total variation distance $\varepsilon/2$ of $\nu_{G, \beta, k}$ on $\Sigma_G(k)$. If no failure occurs then the algorithm must reach a value of λ such that $|n\eta_G(\beta, \lambda) - k| \leq 1/2$. This is because the starting search set Λ_0 contains such a value by Lemma [16](#), and by binary search structure of the algorithm. Note that if there is no failure then $\bar{k}$ is an accurate representation of $n\eta_G(\beta, \lambda)$, so we continue searching in the larger half of the search set when $\bar{k} < k - 1/4$ and so $n\eta_G(\beta, \lambda) \leq \bar{k} + 1/4 < k$. The case that $\bar{k} > k + 1/4$ is similar. The desired total variation distance now follows from the fact that for any λ' , $\nu_{G, \beta, k}$ is precisely $\mu_{G, \beta, \lambda'}$ conditioned on getting a spin assignment of magnetization exactly k . Under the assumption that $\hat{\mu}_{\beta, \lambda} = \mu_{G, \beta, \lambda}$, this means that if the algorithm outputs a state on line (d) during the FOR loop, then the output distribution is precisely $\nu_{G, \beta, k}$. Since we have proved that a state is output on line (d) of some step with probability at least $\varepsilon/2$, this is equivalent to showing that when $\hat{\mu}_{\beta, \lambda} = \mu_{G, \beta, \lambda}$ the output distribution is within total variation distance $\varepsilon/2$ of $\nu_{G, \beta, k}$.

We do not need to assume access to efficient algorithms for sampling from $\mu_{G, \beta, \lambda}$ exactly, we can make do with good approximate samplers given by Theorem [9](#). A standard interpretation of total variation distance is that when each $\hat{\mu}_{\beta, \lambda}$ has total variation distance ξ from $\mu_{G, \beta, \lambda}$, there is a coupling between the measures such that the probability they disagree is at most ξ . Then to prove Theorem [18](#) in general we can add a third failure condition: that any sample from any $\hat{\mu}_{\beta, \lambda}$ disagrees with $\mu_{G, \beta, \lambda}$ under this coupling. We make at most $CN \log n$ calls to any approximate sampling algorithm, so by a union bound this type of failure occurs with probability at most $\varepsilon/2$ when we have the stated total variation distance $\xi \leq \varepsilon/(2CN \log n)$. Now, with the above proof in the case of exact samplers we output a state distributed according to $\nu_{G, \beta, k}$ with probability at least $1 - \varepsilon/2$, and with approximate

samplers as described we get the same output unless the third type of failure occurs. That is, with probability at least $1 - \varepsilon$ we output a state distributed according to $\nu_{G,\beta,k}$. In terms of total variation distance, this is an ε -approximate sampler for $\nu_{G,\beta,k}$ with running time $O(N \log n \cdot T(n, \varepsilon))$ as desired. $\square$

References

- [Bax82] Rodney J. Baxter. *Exactly Solved Models in Statistical Mechanics*. Academic Press, 1982.
- [BJ92] Thang Nguyen Bui and Curt Jones. Finding good approximate vertex and edge partitions is NP-hard. *Information Processing Letters*, 42(3):153–159, 1992.
- [BKW09] Etienne P Bernard, Werner Krauth, and David B Wilson. Event-chain Monte Carlo algorithms for hard-sphere systems. *Physical Review E*, 80(5):056704, 2009.
- [CCM99] N Cancrini, F Cesi, and F Martinelli. The spectral gap for the Kawasaki dynamics at low temperature. *Journal of Statistical Physics*, 95(1):215–271, 1999.
- [CGG⁺16] Jin-Yi Cai, Andreas Galanis, Leslie Ann Goldberg, Heng Guo, Mark Jerrum, Daniel Štefankovič, and Eric Vigoda. #BIS-hardness for 2-spin systems on bipartite bounded degree graphs in the tree non-uniqueness region. *Journal of Computer and System Sciences*, 82(5):690–711, 2016.
- [CGŠVar] Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Sampling colorings and independent sets of random regular bipartite graphs in the non-uniqueness region. *SODA 2022*, to appear.
- [CM00] N Cancrini and F Martinelli. On the spectral gap of Kawasaki dynamics under a mixing condition revisited. *Journal of Mathematical Physics*, 41(3):1391–1423, 2000.
- [Csi16] Péter Csikvári. Extremal regular graphs: the case of the infinite regular tree. *arXiv preprint arXiv:1612.01295*, 2016.
- [Csi17] Péter Csikvári. Lower matching conjecture, and a new proof of Schrijver’s and Gurvits’s theorems. *Journal of the European Mathematical Society*, 19(6):1811–1844, 2017.
- [DF91] Martin Dyer and Alan Frieze. Computing the volume of convex bodies: A case where randomness provably helps. In Béla Bollobás, editor, *Proceedings of Symposia in Applied Mathematics*, volume 44, pages 123–169. American Mathematical Society, Providence, Rhode Island, 1991.
- [DFJ02] Martin Dyer, Alan Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. *SIAM Journal on Computing*, 31(5):1527–1541, 2002.

- [DJPR17] Ewan Davies, Matthew Jenssen, Will Perkins, and Barnaby Roberts. Independent sets, matchings, and occupancy fractions. *J. London Math. Soc.*, 96(1):47–66, August 2017.
- [DM10] Amir Dembo and Andrea Montanari. Ising models on locally tree-like graphs. *The Annals of Applied Probability*, 20(2):565–592, 2010.
- [DMS17] Amir Dembo, Andrea Montanari, and Subhabrata Sen. Extremal cuts of sparse random graphs. *The Annals of Probability*, 45(2):1190–1217, 2017.
- [Dob68] R. L. Dobrushin. Gibbsian random fields for lattice systems with pairwise interactions. *Functional Analysis and its Applications*, 2(4):292–301, October 1968.
- [DP21] Ewan Davies and Will Perkins. Approximately counting independent sets of a given size in bounded-degree graphs. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198, pages 62:1–62:18, 2021.
- [DSW04] Josep Díaz, Maria J Serna, and Nicholas C Wormald. Computation of the bisection width for random d-regular graphs. In *Latin American Symposium on Theoretical Informatics*, pages 49–58. Springer, 2004.
- [DT77] RL Dobrushin and Brunello Tirozzi. The central limit theorem and the problem of equivalence of ensembles. *Communications in Mathematical Physics*, 54(2):173–192, 1977.
- [GJ07] Leslie Ann Goldberg and Mark Jerrum. The complexity of ferromagnetic Ising with local fields. *Combinatorics, Probability and Computing*, 16(1):43–61, 2007.
- [GJ18] Heng Guo and Mark Jerrum. Random cluster dynamics for the Ising model is rapidly mixing. *The Annals of Applied Probability*, 28(2):1292–1313, 2018.
- [GJS74] Michael R Garey, David S Johnson, and Larry Stockmeyer. Some simplified NP-complete problems. In *Proceedings of the sixth annual ACM symposium on Theory of computing*, pages 47–63, 1974.
- [Gri67] Robert B Griffiths. Correlations in Ising ferromagnets. I. *Journal of Mathematical Physics*, 8(3):478–483, 1967.
- [GŠV15] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability for anti-ferromagnetic spin systems in the tree nonuniqueness region. *J. ACM*, 62(6):1–60, December 2015.
- [GŠV16] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability and Computing*, 25(4):500–559, July 2016.
- [GŠVY16] Andreas Galanis, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Ferromagnetic Potts model: Refined #BIS-hardness and related results. *SIAM Journal on Computing*, 45(6):2004–2065, January 2016.

- [HJP20] Tyler Helmuth, Matthew Jenssen, and Will Perkins. Finite-size scaling, phase coexistence, and algorithms for the random cluster model on random graphs. *arXiv preprint arXiv:2006.11580*, 2020.
- [Jan95] Svante Janson. Random regular graphs: asymptotic distributions and contiguity. *Combinatorics, Probability and Computing*, 4(4):369–405, 1995.
- [JKP20] Matthew Jenssen, Peter Keevash, and Will Perkins. Algorithms for #BIS-hard problems on expander graphs. *SIAM Journal on Computing*, 49(4):681–710, 2020.
- [JPPar] Matthew Jenssen, Will Perkins, and Aditya Potukuchi. Approximately counting independent sets in bipartite graphs via graph containers. *SODA 2022*, to appear.
- [JPSS21] Vishesh Jain, Will Perkins, Ashwin Sah, and Mehtaab Sawhney. Approximate counting and sampling via local central limit theorems. *arXiv preprint arXiv:2108.01161*, 2021.
- [JS89] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989.
- [JS93] Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22(5):1087–1116, October 1993.
- [JS97] Mark Jerrum and Alistair Sinclair. The Markov chain Monte Carlo method: An approach to approximate counting and integration. In *Approximation Algorithms for NP-Hard Problems*. PWS Pub. Co, 1997.
- [Kah01] Jeff Kahn. An entropy approach to the hard-core model on bipartite graphs. *Combinatorics, Probability and Computing*, 10(3):219–237, 2001.
- [Kaw66] Kyozi Kawasaki. Diffusion constants near the critical point for time-dependent Ising models. I. *Physical Review*, 145(1):224, 1966.
- [Kaw72] Kyozi Kawasaki. Kinetics of Ising models. *Phase Transitions and Critical Phenomena 2*, pages 443–501, 1972.
- [Kri75] Samuel Krinsky. Bethe approximation to the magnetization is upper bound for Ising ferromagnet. *Physical Review B*, 11(5):1970, 1975.
- [KS68] Douglas G Kelly and Seymour Sherman. General Griffiths’ inequalities on correlations in Ising ferromagnets. *Journal of Mathematical Physics*, 9(3):466–484, 1968.
- [LR69] O. E. Lanford and D. Ruelle. Observables at infinity and states with short range correlations in statistical mechanics. *Communications in Mathematical Physics*, 13(3):194–215, September 1969.
- [LY52] Tsung-Dao Lee and Chen-Ning Yang. Statistical theory of equations of state and phase transitions. II. Lattice gas and Ising model. *Physical Review*, 87(3):410, 1952.

- [LY93] Sheng Lin Lu and Horng-Tzer Yau. Spectral gap and logarithmic Sobolev inequality for Kawasaki and Glauber dynamics. *Communications in Mathematical Physics*, 156(2):399–433, 1993.
- [Lyo89] Russell Lyons. The Ising model and percolation on trees and tree-like graphs. *Communications in Mathematical Physics*, 125(2):337–353, 1989.
- [MP87] Marc Mézard and Giorgio Parisi. Mean-field theory of randomly frustrated systems with finite connectivity. *EPL (Europhysics Letters)*, 3(10):1067, 1987.
- [MRR⁺53] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [MS13] Elchanan Mossel and Allan Sly. Exact thresholds for Ising–Gibbs samplers on general graphs. *The Annals of Probability*, 41(1):294–328, 2013.
- [MS19] Marcus Michelen and Julian Sahasrabudhe. Central limit theorems and the geometry of polynomials. *arXiv preprint*, August 2019.
- [PIG⁺08] Allon G Percus, Gabriel Istrate, Bruno Gonçalves, Robert Z Sumi, and Stefan Boettcher. The peculiar phase structure of random graph bisection. *Journal of Mathematical Physics*, 49(12):125219, 2008.
- [PR20] Han Peters and Guus Regts. Location of zeros for the partition function of the Ising model on bounded degree graphs. *Journal of the London Mathematical Society*, 101(2):765–785, April 2020.
- [Ruo12] Nicholas Ruozzi. The Bethe partition function of log-supermodular graphical models. In *Proceedings of the 25th International Conference on Neural Information Processing Systems-Volume 1*, pages 117–125, 2012.
- [RW94] Robert W. Robinson and Nicholas C. Wormald. Almost all regular graphs are Hamiltonian. *Random Structures & Algorithms*, 5(2):363–374, 1994.
- [RW99] Dana Randall and David Wilson. Sampling spin configurations of an Ising system. In *Symposium on Discrete Algorithms: Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms*, volume 17, pages 959–960, 1999.
- [Sly10] Allan Sly. Computational transition at the uniqueness threshold. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 287–296, Las Vegas, NV, USA, October 2010. IEEE.
- [SS14] Allan Sly and Nike Sun. Counting in two-spin models on d-regular graphs. *Ann. Probab.*, 42(6):2383–2416, November 2014.
- [SSSZ20] Ashwin Sah, Mehtaab Sawhney, David Stoner, and Yufei Zhao. A reverse Sidorenko inequality. *Inventiones mathematicae*, 221(2):665–711, 2020.

- [SST14] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *Journal of Statistical Physics*, 155(4):666–686, 2014.
- [ŠVV09] Daniel Štefankovič, Santosh Vempala, and Eric Vigoda. Adaptive simulated annealing: A near-optimal connection between sampling and counting. *Journal of the ACM*, 56(3):18:1–18:36, May 2009.
- [Tho71] Colin J Thompson. Upper bounds for Ising model correlation functions. *Communications in Mathematical Physics*, 24(1):61–66, 1971.
- [Wei06] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing - STOC '06*, page 140, Seattle, WA, USA, 2006. ACM Press.
- [Yau96] Horng-Tzer Yau. Logarithmic Sobolev inequality for lattice gases with mixing conditions. *Communications in Mathematical Physics*, 181(2):367–408, 1996.
- [ZB10] Lenka Zdeborová and Stefan Boettcher. A conjecture on the maximum cut and bisection width in random regular graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(02):P02020, 2010.
- [Zha17] Yufei Zhao. Extremal regular graphs: independent sets and graph homomorphisms. *The American Mathematical Monthly*, 124(9):827–843, 2017.

A Proof of Proposition 7

First we prove the variance bound.

Proof. Fix Δ, β, λ satisfying the assumptions of the proposition, and let $G \in \mathcal{G}_\Delta$ be a graph on n vertices. For the lower bound, fix an independent set J in G of size at least $n/(\Delta + 1)$, let $U = V(G) \setminus J$, and let σ_U, σ_J be σ restricted to U and J respectively. Then $\text{var}(\mathbf{X}) \geq \langle \text{var}(\mathbf{X}|\sigma_U) \rangle_{G, \beta, \lambda}$ by the law of total variance. But conditioned on σ_U , the spins of the vertices in J are independent, and moreover, the conditional marginal probability of each spin taking $+$ is bounded away from 0 and 1 independently of n . That is, conditioned on σ_U the indicator random variable $\mathbf{X}_v = \mathbf{1}_{\sigma_v=1}$ for v in J are mutually independent Bernoulli random variables with

$$\text{var}(\mathbf{X}_v|\tau) = \frac{1}{(\lambda e^{\frac{\beta}{2}Y_{v,\tau}} + \lambda^{-1}e^{-\frac{\beta}{2}Y_{v,\tau}})^2},$$

where $Y_{v,\tau}$ is the magnetization under σ_U of the neighbors of v (so $Y_{v,\tau} = 2j - \deg(v)$ when v has j neighbors with spin $+$ under σ_U). Since we have $\beta \geq 0$ and $\lambda \geq 1$, it is straightforward to prove that $\text{var}(\mathbf{X}_v|\tau)$ is minimized when $Y_{v,\tau} = \deg(v)$, and is a decreasing function of $\deg(v)$. We conclude that

$$\text{var}(\mathbf{X}|\sigma_U) = \sum_{v \in J} \text{var}(\mathbf{X}_v|\sigma_U) \geq |J| \frac{1}{(\lambda e^{\frac{\beta\Delta}{2}} + \lambda^{-1}e^{-\frac{\beta\Delta}{2}})^2}.$$

For the upper bound, we proceed as in [JPSS21, Lemma 3.2]. Let $\hat{Z}_G(\beta, t) = \lambda^n Z_G(\beta, \lambda)$ with $t = \sqrt{\lambda}$. That is,

$$\hat{Z}_G(\beta, t) = \sum_{\sigma \in \Sigma_G} e^{\frac{\beta}{2} \sum_{(u,v) \in E(G)} \sigma_u \sigma_v} \cdot t^{X(\sigma)}.$$

This defines the same Ising model, but will be slightly easier to work with since with β fixed, the partition function $\hat{Z}_G(\beta, t)$ is a polynomial in t of degree n . Let $\xi_1, \dots, \xi_n$ denote its complex roots. The assumptions of Proposition 4, along with Theorems 4 and 5, imply that there exists some $\delta > 0$, depending only on Δ , β , and λ so that for each j ,

$$|\xi_j - t| > \delta \tag{27}$$

with $|\cdot|$ denoting distance in the complex plane. We write

$$\hat{Z}_G(\beta, t) = \prod_{j=1}^n (1 - t/\xi_j)$$

and then

$$\begin{aligned} \text{var}(\mathbf{X}) &= t^2 \frac{\partial^2 \log \hat{Z}_G(\beta, t)}{\partial t^2} + t \frac{\partial \log \hat{Z}_G(\beta, t)}{\partial t} \\ &= -t \sum_{j=1}^n \frac{1}{\xi_j (1 - t/\xi_j)^2}. \end{aligned}$$

From Theorem 4 we have $|\xi_j| = 1$ for all j and from (27) we have $|1 - t/\xi_j|^{-2} \leq \delta^{-2}$, giving $\text{var}(\mathbf{X}) \leq \frac{n\sqrt{\lambda}}{\delta^2}$. $\square$

Next we prove the local central limit theorem. As remarked above, this essentially follows from the method of [DT77], as it is straightforward to generalize their proof from $\mathbb{Z}^d$ to general bounded-degree graphs.

Proof. Again fix Δ, β, λ satisfying the assumptions of the proposition, and let $G \in \mathcal{G}_\Delta$ be a graph on n vertices. We let $\langle \cdot \rangle$ denote expectation with respect to $\mu_{G, \beta, \lambda}$.

Let $\phi_{\mathbf{X}}(t) = \langle e^{it\mathbf{X}} \rangle$ denote the characteristic function of $\mathbf{X}$, and let $\kappa^2 = \langle (\mathbf{X} - \langle \mathbf{X} \rangle)^2 \rangle$ be the variance of $\mathbf{X}$. Let $\bar{\mathbf{X}} = (\mathbf{X} - \langle \mathbf{X} \rangle)/\kappa$, and let $\phi_{\bar{\mathbf{X}}}(t)$ be the characteristic function of $\bar{\mathbf{X}}$. Let $\mathcal{L}$ denote the lattice $\langle \mathbf{X} \rangle + \mathbb{Z}/\kappa$. Let $\mathcal{N}(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$. We want to show that

$$\sup_{x \in \mathcal{L}} \left| \mu_{G, \beta, \lambda}(\bar{\mathbf{X}} = x) - \frac{1}{\kappa} \mathcal{N}(x) \right| = o(n^{-1/2}),$$

or

$$\sup_{x \in \mathcal{L}} \left| \kappa \mu_{G, \beta, \lambda}(\bar{\mathbf{X}} = x) - \mathcal{N}(x) \right| = o(1),$$

since $\kappa = \Theta(\sqrt{n})$. Using Fourier inversion we have

$$\begin{aligned}
& 2\pi \sup_{x \in \mathcal{L}} |\kappa \mu_{G, \beta, \lambda}(\bar{\mathbf{X}} = x) - \mathcal{N}(x)| \\
&= \sup_{x \in \mathcal{L}} \left| \int_{-\pi\kappa}^{\pi\kappa} \phi_{\bar{\mathbf{X}}}(t) e^{-itx} dt - \int_{-\infty}^{\infty} e^{-t^2/2 - itx} dt \right| \\
&\leq \int_{-\pi\kappa}^{\pi\kappa} |\phi_{\bar{\mathbf{X}}}(t) - e^{-t^2/2}| dt + \int_{|t| > \pi\kappa} e^{-t^2/2} dt \\
&\leq \int_{-K}^K |\phi_{\bar{\mathbf{X}}}(t) - e^{-t^2/2}| dt + \int_{|t| \geq K} e^{-t^2/2} dt + \int_{|t| \geq K} |\phi_{\bar{\mathbf{X}}}(t)| dt \\
&=: A_1 + A_2 + A_3.
\end{aligned}$$

It is enough to show that for every $\varepsilon > 0$, there is n large enough so that $A_1 + A_2 + A_3 \leq \varepsilon$. To do this we will choose K large enough as a function of ε .

Because $\lim_{n \rightarrow \infty} \phi_{\bar{\mathbf{X}}}(t) = e^{-t^2/2}$ for every fixed t (by the central limit theorem of Theorem 6 and the zero-freeness results supplied by the conditions of the proposition and Theorems 4 and 5), applying the bounded convergence theorem gives that $A_1 \rightarrow 0$ as $n \rightarrow \infty$, so we can choose n large enough to guarantee $A_1 < \varepsilon/3$.

We can pick K large enough to ensure $A_2 < \varepsilon/3$. Similarly, if we can show that there exists some $c > 0$ so that for every $t \in \mathbb{R}$, $|\phi_{\bar{\mathbf{X}}}(t)| \leq e^{-ct^2}$, then we can choose K large enough to ensure $A_3 < \varepsilon/3$, as desired.

To show this, it is enough to prove that $|\langle e^{it\mathbf{X}} \rangle| \leq e^{-cnt^2}$. To do this, we proceed as in the proof of the variance lower bound as above. Let $J \subset V(G)$ be an independent set of size at least $n/(\Delta + 1)$ and let $U = V(G) \setminus J$. We will again use the fact that conditioned on σ_U , the spins of the vertices in J are independent with bounded marginals. Let $\mathbf{X} = \mathbf{X}_U + \mathbf{X}_J$ where $\mathbf{X}_U, \mathbf{X}_J$ are the number of $+$ spins in the vertices in U and J respectively. Then

$$\begin{aligned}
|\langle e^{it\mathbf{X}} \rangle| &= |\langle e^{it\mathbf{X}_U} e^{it\mathbf{X}_J} \rangle| \\
&\leq \langle |e^{it\mathbf{X}_U}| \langle e^{it\mathbf{X}_J} | \sigma_U \rangle \rangle \\
&\leq \langle |e^{it\mathbf{X}_J}| | \sigma_U \rangle \rangle \\
&\leq \max_{\tau_U} |\langle e^{it\mathbf{X}_J} | \sigma_U = \tau_U \rangle|.
\end{aligned}$$

We will show that for all spin assignments τ_U to U , $|\langle e^{it\mathbf{X}_J} | \sigma_U = \tau_U \rangle| \leq e^{-cnt^2}$. This follows from the independence and bounded marginals properties of the spins in J conditioned on σ_U . In particular, there exists $c' > 0$ so that $|\langle e^{it(1+\sigma_v)/2} | \sigma_U = \tau_U \rangle| \leq e^{-c't^2}$ over all $v \in J$ and all choices of τ_U . Then by independence, $|\langle e^{it\mathbf{X}_J} | \sigma_U = \tau_U \rangle| \leq e^{-c'nt^2/(\Delta+1)}$, and the claim follows by taking $c = c'/(\Delta + 1)$. $\square$

B Approximate counting via sampling

In this section we fix $\Delta \geq 3$, $\beta \in (\beta_c(\Delta), 1)$, $\eta \in [0, 1)$, an n -vertex graph $G \in \mathcal{G}_\Delta$, and k such that $0 \leq k \leq \eta n$ and $n \equiv k \pmod{2}$. Since G and k remains fixed we freely drop them

from notation such as $Z_G^{\text{fix}}(\beta, k)$ and $\nu_{G, \beta, k}$. Reducing approximate counting to approximate sampling is a well-studied area, and a standard simulated annealing approach lets us do this for $Z_G^{\text{fix}}(\beta)$ and ν_β . For simplicity, we present only a basic form of the argument. More sophisticated cooling schedules exist that could be used to improve the running time of the reduction. The key observations behind this application of simulated annealing are that $Z^{\text{fix}}(1) = \binom{n}{(n+k)/2}$, and for any valid magnetization k and any parameters β, β' we have

$$\frac{Z^{\text{fix}}(\beta')}{Z^{\text{fix}}(\beta)} = \left\langle e^{\frac{\beta' - \beta}{2} \delta(\sigma)} \right\rangle_{\beta, k}, \quad (28)$$

where $\delta(\sigma) = \sum_{uv \in E(G)} \sigma_u \sigma_v$ and $\langle \cdot \rangle_{\beta, k}$ represents an expectation over $\nu_{G, \beta, k}$. These facts follow straightforwardly from the definition of $Z^{\text{fix}}(\beta)$. Given a *cooling schedule*

$$1 = \beta_0 < \beta_1 < \dots < \beta_\ell = \beta,$$

and the independent random variables $R_i = e^{\frac{\beta_{i+1} - \beta_i}{2} \delta(\sigma_i)}$ where $\sigma_i \sim \nu_{\beta_i}$, we have

$$Z^{\text{fix}}(\beta) = \binom{n}{(n+k)/2} \prod_{i=0}^{\ell-1} \langle R_i \rangle_{\beta, k}.$$

Moreover, by (28) we can estimate $\langle R_i \rangle_{\beta, k}$ by taking S repeated independent samples from a ξ -approximate sampler $\hat{\nu}_i$ for ν_{β_i} , and setting $\hat{R}_i$ to be the mean of the S samples. For the cooling schedule to be useful in an FPRAS, we need to determine suitable S and ξ so that

$$\mathbb{P} \left((1 - \varepsilon) Z_G^{\text{fix}}(\beta, k) \leq \binom{n}{(n+k)/2} \prod_{i=0}^{\ell} \hat{R}_i \leq (1 + \varepsilon) Z_G^{\text{fix}}(\beta, k) \right) \geq 3/4. \quad (29)$$

An important consideration for a cooling schedule is the quantity

$$\frac{\langle R_i^2 \rangle_{\beta, k}}{\langle R_i \rangle_{\beta, k}^2} = \frac{Z^{\text{fix}}(2\beta_{i+1} - \beta_i) Z^{\text{fix}}(\beta_i)}{Z^{\text{fix}}(\beta_{i+1})^2},$$

and when this is bounded above by a constant B for all $0 \leq i \leq \ell$ we say the schedule is *B-Chebyshev*. A particularly simple cooling schedule has $\beta_i = i \log(1 + 1/n)$ for $i < \ell$, so that for some length $\ell = \Theta(n)$ we have a schedule where $\beta_{i+1} - \beta_i \leq \log(1 + 1/n)$. Then from the fact that when $r \geq 1$ we have $Z^{\text{fix}}(\beta + r) \leq r^{\Delta n/2} Z^{\text{fix}}(\beta)$, it is straightforward to show that $1 \leq R_i \leq e^{\Delta/2}$ and

$$\frac{\langle R_i^2 \rangle_{\beta, k}}{\langle R_i \rangle_{\beta, k}^2} \leq (1 + 1/n)^{\Delta n/2} \leq e^{\Delta/2},$$

meaning that the simple cooling schedule is $e^{\Delta/2}$ -Chebyshev. A second-moment argument, given for example in [SVV09, Section 2] and in [JS97] (but originating in [DF91]), gives that some $S = \Omega(\ell/\varepsilon^2)$ and $\xi = O(\varepsilon/\ell)$ suffice for (29) (where the implied constants can depend on Δ). By the algorithm from Section 5, we can obtain such ξ -approximate samples in time polynomial in n and $\log(1/\varepsilon)$. The total running time for the approximate counting algorithm is therefore polynomial in n and $1/\varepsilon$, meaning we have the desired FPRAS.

This simple cooling schedule works in great generality, but there are a few well-known ways of improving the efficiency of the method. The first is to use a more complex (in fact, *adaptive*) but shorter schedule, as given in [ŠVV09]. A second improvement makes use of *warm starts*, which is essentially a strategy of reusing randomness when a Markov chain is the basis of the approximate samplers $\hat{\nu}_i$ which are used to obtain $\hat{R}_i$. These improvements are combined in [ŠVV09, Section 7].